

Spike2 for Windows

Version 7

Copyright © Cambridge Electronic Design Limited 1988-2012

Neither the whole nor any part of the information contained in, or the product described in, this guide may be adapted or reproduced in any material form except with the prior written approval of Cambridge Electronic Design Limited.

Version 7.00	June 2009
Version 7.01	July 2009
Version 7.02	November 2009
Version 7.03	February 2010
Version 7.04	June 2010
Version 7.05	Oct 2010
Version 7.06	Jan 2011
Version 7.07	May 2011
Version 7.08	October 2011
Version 7.09	January 2012
Version 7.10	May 2012
Version 7.11	Dec 2012

Published by:

Cambridge Electronic Design Limited
Science Park
Milton Road
Cambridge
CB4 0FE
UK

Telephone: Cambridge (01223) 420186
International +44 1223 420186
Fax: Cambridge (01223) 420488
International +44 1223 420488
Email: info@ced.co.uk
Home page: www.ced.co.uk

Trademarks and Tradenames used in this guide are acknowledged to be the Trademarks and Tradenames of their respective Companies and Corporations.

Table of Contents

Spike2 for Windows version 7	1-1
New features in version 7.....	1-1
File icons.....	1-2
Licence information	1-2
Direct access to the raw data.....	1-3
Versions of Spike2.....	1-3
Using this manual	1-3
Hardware required.....	1-3
Installing, updating and removing Spike2.....	1-4
Getting started	2-1
Open Spike2.....	2-1
Scroll bar and associated buttons	2-2
Vertical cursors	2-2
Active cursor searches	2-3
Step to next, previous event	2-3
Zoom using the mouse	2-3
Drag axes	2-4
X axis dialog	2-4
Y axis dialog	2-5
Channel draw mode	2-5
Show/hide channels.....	2-5
Cursor Y values.....	2-6
Cursor regions.....	2-6
Channel order.....	2-7
Channel size	2-7
Channel overdraw	2-8
Settings dialog.....	2-8
Create result view.....	2-9
Process data.....	2-10
Draw mode.....	2-10
Cursors	2-10
Measurements	2-10
Create with a script	2-12
General information.....	3-1
Channel lists.....	3-1
Dialog expressions	3-1
Data view keyboard shortcuts	3-3
Text view keyboard shortcuts	3-3
The Spike2 command line.....	3-6
Shell extensions	3-6
Sampling data	4-1
Sampling configuration.....	4-1
Channels.....	4-1
Talkers	4-10
Resolution	4-12
Sampling mode	4-17
Sequencer.....	4-19
Arbitrary waveform output	4-20
Play waveform	4-22
Automation	4-24
Outputs.....	4-26
Process dialog for a new file	4-28
Sample control toolbar	4-28

High sampling rates4-29
The Sample Status bar4-29
Saving configurations4-30

Data output during sampling 5-1

Sequencer control panel.....5-1
Creating sequences5-2
Sequencer speed.....5-2
Marker channel and digital input conflict5-3
Digital input and output5-4
The graphical editor5-5
Graphical editing5-8
Graphical palette5-9
The text editor5-14
Instructions5-17
Instruction format5-18
Expressions5-18
Variables5-19
Constants5-20
The SET, SCLK and SDAC directives5-20
Table of values.....5-21
Include files5-22
Changing the sequence during sampling5-23
Sequencer instruction reference.....5-24
Digital I/O.....5-24
DAC outputs5-27
Cosine output control instructions5-29
General control5-35
Variable arithmetic5-37
Variable logic5-40
Table access.....5-40
Access to data capture.....5-42
Randomisation5-43
Arbitrary waveform output5-44
Compatibility with previous versions5-46

File menu 6-1

New File6-1
Open6-2
Import6-2
Global Resources6-2
Resource Files.....6-3
Close6-4
Revert To Saved6-4
Save and Save As.....6-4
Export As6-4
Load and Save Configuration6-6
Page Setup6-6
Page Headers6-7
Print Preview6-8
Print Visible, Print and Print Selection6-8
Print Screen.....6-9
Exit6-9
Send Mail.....6-10
Read-only files.....6-10

Edit menu	7-1
Undo and Redo	7-1
Cut, Copy, Paste, Delete and Clear	7-1
Copy Spreadsheet.....	7-2
Copy As Text	7-2
Select All.....	7-6
Find, Find Again and Find Last	7-7
Replace.....	7-8
Edit toolbar.....	7-8
Auto Format	7-9
Toggle Comments	7-9
Auto Complete	7-10
Preferences	7-11
View menu	8-1
Enlarge View and Reduce View	8-1
X, Y Axis Range	8-1
Standard Display	8-3
Show/Hide Channel	8-3
Info.....	8-4
File Information	8-4
Channel Information	8-4
Options.....	8-5
Trigger/Overdraw.....	8-5
Display Trigger	8-6
Overdraw List	8-8
Overdraw 3D.....	8-8
Clear List.....	8-9
Channel Draw Mode	8-10
XY Draw Mode.....	8-14
Vertical Marker	8-15
Multimedia files	8-16
Font	8-17
Change Colours.....	8-18
Use Colour and Use Black And White.....	8-20
Sonogram Colours	8-20
ReRun	8-21
Pen Width.....	8-22
Channel Image	8-22
Show Gutter	8-24
Show Line Numbers.....	8-24
Folding	8-24
Spike Monitor	8-24
Analysis menu.....	9-1
New Result View	9-1
Interval histogram	9-2
Peri-stimulus time histogram	9-3
Event correlations	9-4
Phase histogram	9-6
Waveform average	9-7
Power spectrum.....	9-8
Waveform correlation	9-11
Process settings	9-11
Process	9-12
Process command with a new file.....	9-13
Measurements to XY views	9-14

Measurements to a data channel	9-20
Fit Data	9-21
Memory buffer.....	9-26
Virtual Channels	9-28
Duplicate Channels.....	9-34
Save channel	9-35
Delete channel	9-35
Calibrate.....	9-35
Channel process	9-38
Marker Filter.....	9-40
Set Marker Codes	9-41
New WaveMark.....	9-42
New n-trode	9-42
Edit WaveMark.....	9-42
Digital filters.....	9-42

Window menu 10-1

Duplicate window.....	10-1
Hide	10-1
Show	10-1
Tile Horizontally.....	10-1
Tile Vertically.....	10-1
Cascade.....	10-1
Arrange Icons	10-2
Close All.....	10-2
Windows.....	10-2

Cursor menu 11-1

New Cursor.....	11-1
Delete, Fetch, Move To, Display All.....	11-1
Position Cursor	11-1
Label mode	11-2
Set Label.....	11-2
Renumber.....	11-2
Active cursors	11-3
Active mode.....	11-3
Search Right, Search Left.....	11-5
Display y values.....	11-6
Cursor regions.....	11-7
Horizontal cursors.....	11-9
New Horizontal.....	11-9
Delete Horizontal.....	11-9
Fetch Horizontal	11-9
Move To Horizontal	11-9
Position Horizontal	11-9
Display all Horizontal.....	11-9
Horizontal Label mode	11-10
Renumber Horizontal.....	11-10
Cursor context commands	11-10

Sample menu 12-1

Sampling configuration.....	12-1
Clear configuration	12-1
Sample Bar	12-1
Create a TextMark	12-1
Sampling Notes.....	12-2

Talkers	12-2
Offline waveform output.....	12-3
Sampling controls	12-4
Sequencer controls	12-4
Change Output Sequence	12-4
Conditioner Settings.....	12-4
Discriminator Configuration	12-4
Script menu	13-1
Compile and Run Script	13-1
Evaluate	13-1
Turn Recording On/Off.....	13-1
Debug bar.....	13-1
Convert DOS Script	13-1
Script Bar	13-2
Help menu	14-1
Using help.....	14-1
Tip of the Day	14-1
View Web site.....	14-1
Getting started.....	14-1
Other sources of help	14-1
Spike shapes	15-1
Spike detection.....	15-1
Sampling parameters.....	15-2
Template matched signal	15-2
On-line template setup	15-3
Selecting the area for a template	15-4
The template area.....	15-5
Merging templates.....	15-5
Manual template creation.....	15-5
Toolbar controls	15-6
Horizontal cursors.....	15-7
Multiple traces	15-7
Template formation.....	15-8
Template settings dialog	15-9
Off-line template formation	15-11
Off-line template editing.....	15-13
Collision Analysis Mode.....	15-15
On-line template monitoring.....	15-17
Load/Save templates	15-18
Template storage	15-19
Spike Monitor	15-20
Getting started with spike shapes and templates	15-22
Creating on-line templates with cluster analysis.....	15-31
Clustering	16-1
Principal Component Analysis.....	16-2
Cluster on measurements	16-4
Cluster on template correlations.....	16-6
Cluster on template errors	16-7
The Clustering dialog.....	16-8
Menu commands	16-11
File menu	16-11
Edit menu	16-12

- View menu.....16-14
- Use Time as Z.....16-15
- Cluster menu.....16-16
- Time range control.....16-19
- Online clustering.....16-20
- Getting started with clustering.....16-20
- K Means algorithm.....16-21
- Normal Mixtures algorithm.....16-23
- Mahalanobis distance.....16-23

- Digital filtering..... 17-1**
 - Digital filter dialog.....17-2
 - Filter bank.....17-3
 - FIR filter details.....17-4
 - FIR Filter types.....17-5
 - IIR filter details.....17-6
 - FIR filters technical information.....17-9
 - FIRMake() script command.....17-11
 - FIR filter examples.....17-12
 - Hilbert transformer.....17-18

- Programmable signal conditioners..... 18-1**
 - What a signal conditioner does.....18-1
 - Serial ports.....18-1
 - Control panel.....18-1
 - Setting the channel gain and offset.....18-3
 - Conditioner connections.....18-3

- 1401-18 programmable discriminator..... 19-1**
 - What the 1401-18 does.....19-1
 - Signal route.....19-1
 - Discrimination mode.....19-2
 - View data.....19-2
 - Monitor channel.....19-3
 - Digital input connections.....19-4
 - Electrical information.....19-4

- Utility programs..... 20-1**
 - The SonFix data recovery utility.....20-1
 - Try1401 test program.....20-5

- Multimedia recording..... 21-1**
 - System requirements.....21-1
 - Getting started.....21-2
 - Video Device Properties.....21-2
 - Audio Device Properties.....21-3
 - Video Capture.....21-3
 - Set Slow Frame Rate.....21-3
 - Use Slow Frame Rate.....21-3
 - Configuration.....21-4
 - View menu.....21-7
 - File menu.....21-7
 - Recording data.....21-8
 - Timing and AVI files.....21-8
 - The avicomp application.....21-9

Spike2 for Windows version 7

Introduction With Spike2 version 7 and a CED 1401 family interface (not standard 1401) you can capture and analyse waveform, event and marker data and output precisely timed pulses and voltages using the familiar and easy to use Windows environment. If you have a standard 1401 you can sample data with Spike2 version 3 (also on the installation CD).

You can arrange the windows to display the data within them to best advantage and cut and paste the results to other applications. Alternatively, you can obtain printer hard copy directly from the application. When you close a data file, Spike2 saves the screen format and channel display settings. When you open a file, Spike2 restores the configuration, so it is easy to resume work where you stopped in a previous session.

You can analyse sections of data by reading off values at and between cursors, or by applying the built-in functions, for example waveform averaging, digital filtering, spike detection, histogram formation and power spectra. More ambitious users can automate both data capture and analysis with scripts and the output sequencer. The script language is described in *The Spike2 script language* manual. This manual describes the general interactive use of Spike2. You will find installation information on page 4.

New features in version 7 We have tried very hard to keep version 7 of Spike2 compatible with version 6. It reads data files from all previous versions. Resource files are mostly compatible; some resource formats have changed to support new features. Scripts that ran with version 6 should work unchanged with version 7. New features in version 7 include:

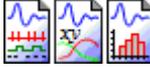
- Data file size can optionally extend to 1 TB (was 2 GB). This is a file format change and such files will be read-only in version 6 and unreadable by previous versions. There is a new version of the SonFix application that knows about big files.
- When exporting you can use big file format or the previous one with a 2 GB limit.
- You can draw event and marker data as Vertical Markers over all the channels to mark points of interest. You can use these as fixed vertical cursors.
- There are new and enhanced drawing modes and you can display the x axis in ms.
- TextMark data can be drawn as Vertical Markers with the text displayed vertically.
- When you open a big file, jumping into the middle of the file is much faster.
- The output sequencer allows sequences of 8191 instructions (was 1023) and up to 256 variables (was 64) and supports constant values and `#include` files.
- The sequencer has new instructions and enhancements to online text displays.
- The clustering display is faster and has new drawing modes.
- You can change to a different output sequence whilst sampling data.
- The current sampling configuration can be copied to the clipboard or Log.
- The sampling configuration dialog is now resizable to cope with long channel lists.
- The Sample Status bar displays remaining percentage of sample time and disk space.
- Scripts can now access user mouse movements and button clicks.
- The script language can now `resize` arrays.
- You can break out of script loops with the `break` and `continue` keywords.
- You can initialise variables in `var` statements with non-constant expressions.
- Script debugger now has a watch window.
- Text views can have line numbers and you can limit the number of lines in a view.
- The new `LinPred()` script command can be used to predict data.
- Virtual channels have comparison operators, `Min()`, `Max()` and power functions.
- Talker interface adds external data to Spike2 file during sampling.

There is a full list of new features, bug fixes and changes in the Revision History in the on-line Help. Licensed users of Spike2 version 7 can download updates of version 7 from our Web site www.ced.co.uk as they become available.

File icons



The various file types in Spike2 have icons so that you can easily recognise them in directory listings. The icon to the left is the Spike2 application icon that you double-click to launch Spike2.



The icon on the left is for Spike2 data files. The icon on the right is for saved result views. The central icon is for an XY file. If you double-click one of these it will launch the Spike2 application (if it is not already running) and open the file.



Spike2 can output sequences of pulses, sine waves and voltage levels as it samples data. Output sequence files have this icon.



This icon is for Spike2 script files. A Spike2 script can automate data capture and analysis operations and extend the capabilities of the Spike2 program.



The icon on the left is for CED configuration files; these hold all the information needed to sample a new data file. The icon on the right is for CED resource files; these are usually associated with a data file and hold things like axis display ranges.

Licence information

CED software is protected by both United Kingdom Copyright Law and International Treaty provisions. Unless you have purchased multiple licences, you are licensed to run one copy of the software. Each copy of the software is identified by a serial number (located in the Help menu About Spike2... dialog box). You may make archival copies of the software for the sole purpose of back up in case of damage to the original. You may install the software on more than one computer as long as there is **No Possibility** of it being used at one location while it is being used at another. If multiple simultaneous use is possible, you must purchase additional software licences.

Additional software licences

The original licensee of a CED software product can purchase additional licences to run multiple copies of the same software. CED supplies an additional manual set with each licence. CED does not supply additional software media. As these additional licences are at a substantially reduced price, there are limitations on their use:

1. The additional licences cannot be separated from the original software and are recorded at CED in the name of the original licensee.
2. All support for the software is expected to be through one nominated person, usually the original licensee.
3. The additional licensed copies are expected to be used on the same site and in the same building/laboratory and by people working within the same group.
4. When upgrades to the software become available that require payment, both the original licence and the additional licences must be upgraded together. If the upgrade price is date dependent, the date used is the date of purchase of the original licence. If some or all of the additional licences are no longer required, you can cancel the unwanted additional licences before the upgrade.
5. If you are the user of an additional licence and circumstances change such that you no longer meet the conditions for use of an additional licence, you may no longer use the software. In this case, with the agreement of the original licensee, it may be possible for you to purchase a full licence at a price that takes into account any monies paid for the additional licence. Contact CED to discuss your circumstances.
6. If you hold the original licence and you move, all licences are presumed to move with you unless you notify us that the software should be registered to someone else.

Direct access to the raw data

Some users may wish to write their own applications that manipulate the Spike2 data files directly. A C library *Spike2: Son data storage library* is available from CED together with documentation sufficient for an experienced C programmer to use it. The library documentation is also available as a pdf file on the Spike2 distribution CD. To install it, select Custom Install and check the Additional documentation box.

Versions of Spike2

This manual describes Spike2 for Windows version 7. The following versions of “Spike2 for ...” exist, listed in more or less chronological order of release:

DOS	Data and output sequencer files from this are compatible with all other versions of Spike2. The script language is different. Newer versions include utilities to convert scripts (as far as is possible).
Macintosh 68k	Data files, scripts and output sequences are compatible with the PowerPC and Windows versions.
Windows version 2	This version will operate with Windows 3.1 and 3.11.
PowerPC Macintosh	This was equivalent to Spike2 for Windows version 2.
Windows version 3	This version supports the standard 1401.
Windows versions 4-6	The previous versions of Spike2.
Windows version 7	The version described in this manual.

Using this manual

The first section of this manual introduces you to Spike2 by suggesting a few tasks you might undertake to get started with the system. We have supplied example data files for you to experiment with; there is no need to have your own data available at this time.

The second section explains the use of the menu commands. We suggest that you work through as much (or as little) of the familiarisation as you feel you need, then dip into the menu section as required for more detailed information.

We do not explain standard procedures, for example clicking and dragging, using menu short-cut keys or using a file open dialog; we expect that you are already familiar with them. We use standard Windows idioms wherever possible so that you feel at home and have a consistent interface to work with.

Once you are familiar with the program, you may wish to investigate the script language so you can automate your data capture and analysis. This is described in the separate reference manual *The Spike2 script language*. The on-line Help system duplicates the information in the manuals and is often the fastest way to look-up a topic. It is usually more up-to-date as the on-line Help is revised with each program revision.

The *Spike2 Training Course Manual* covers selected topics in more detail. It is particularly useful for script authors as the approach is much more descriptive than *The Spike2 script language* reference manual.

Hardware required

The minimum supported hardware for Spike2 version 7 is a 486 computer running Windows XP SP2 with 1 GB of memory. It may run on lower specification systems, but we do not recommend or support this. The more powerful the processor and the more memory your system has, the better Spike2 runs.

To sample data, you need a CED Power1401, Micro1401 or 1401*plus*. The *Owners Handbook* that came with your 1401 has hardware installation instructions. In this manual, Micro1401 is used to describe both the original micro1401 and the Micro1401

mk II and -3 unless the full name is given. Wherever micro1401 is used, it refers to the original micro1401, not the mk II or -3.

Installation The Spike2 installation disk installs the Spike2 program, plus all required drivers for your 1401 and a 1401 test program to verify correct system operation. You do not need to use any other CED installation disk. Your installation disk is serialised to personalise it to you. Please do not allow others to install unlicensed copies of Spike2.

Just put the CD-ROM in the drive and it will start the installation. You can also run the installation manually by opening the folder `Spike7` on the CD-ROM, then open the `disk1` folder and run `setup.exe`.

During installation You must select a suitable drive and folder for Spike2 and personalise your copy with your name and organisation. You can have versions 2 to 7 on the same system as long as they are in different folders. If you have a previous version on your system, install to a different folder. The installation program copies the Spike2 program plus help, demonstration, example and tutorial files. It also copies and installs all required 1401 support (device drivers and control panels). In rare cases you may need to install the drivers manually; the installation program will tell you if this is the case and give you detailed instructions. Your system may require a restart after installation to get all drivers up to date.

Custom install To install without 1401 support, or to copy additional information, which includes the latest copies of the Spike2 manuals, the SON filing system documentation, notes on the Fast Fourier Transform and some technical notes on interfacing to Spike2 sampling, choose **Custom** installation.

After Installation If you are new to Spike2, please work through the *Getting Started* tutorial in the *Spike2 for Windows* manual. Where you go next depends on your requirements. The *Spike2 Training Course Manual* is more descriptive than the other manuals, which are organised as reference material. However, it covers all versions of Spike2 and you will occasionally need to refer to the other manuals for version 7 specific details. The on-line Help in Spike2 has a lot of information; if in doubt use the `F1` key for Help.

Updating Spike2 You can update your copy of Spike2 version 7 to the latest version 7 release from our Web site: www.ced.co.uk. You can only update a correctly installed copy of Spike2 version 7. There are full instructions for downloading the update on the Web site.

Once you have downloaded the Spike2 update, you will find that the update program is very similar to the original installation, except that you must already have a properly installed copy of Spike2 for Windows version 7 on your computer.

Updates will include both bug fixes and new features. We will notify you by email (if we know your email address) of new releases. You can also register for this service on our web site. To stop emails, reply to them and ask to be removed from the list.

Removing Spike2 To remove Spike2: open the system Control Panel, select **Add/Remove Programs**, select **CED Spike2 for Windows version 7** and click **Remove**. This removes files installed with Spike2; you will not lose files you created.

Getting started

Introduction

In this section you will open a Spike2 data file, manipulate the contents and familiarise yourself with the basic display controls. Instructions that you must follow to keep in step with the text are in **bold** type with a pointing finger. Explanations are in normal type.

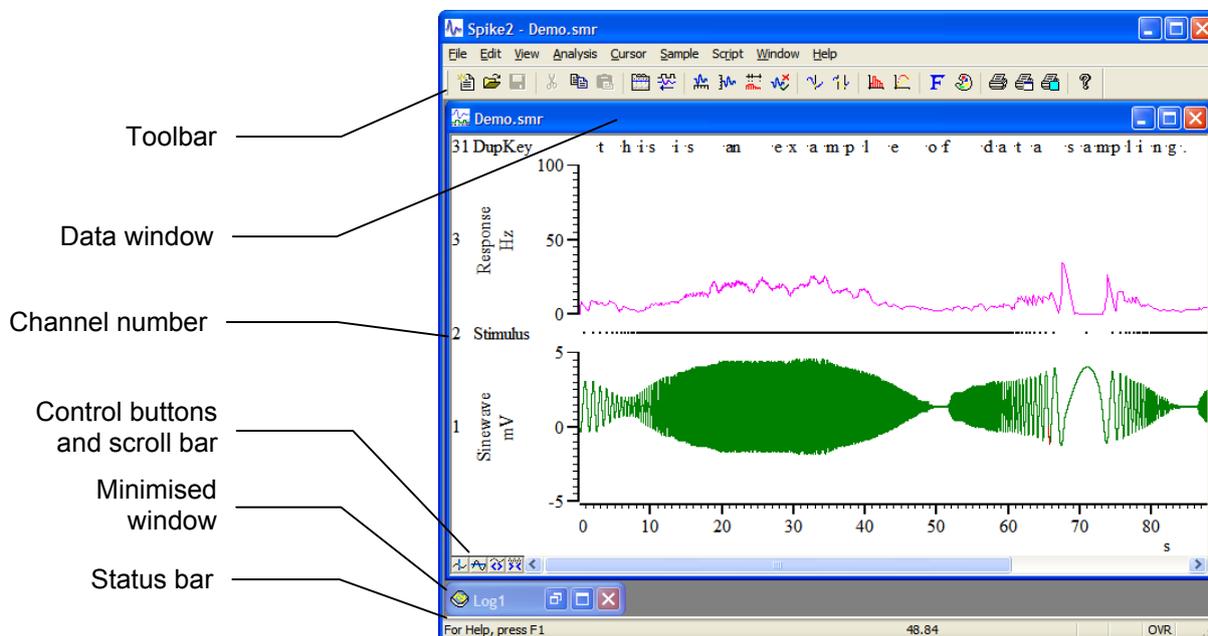
Basic operations

The first task is to become familiar with the basic operations that are always needed to manipulate Spike2 files. We use a sample file called `demo.smr`. The Help menu Index option *Getting started* topic duplicates this chapter. Once you have started Spike2 you may prefer to follow the remainder of the chapter from the on-screen help.



From the Start button choose Programs, then Spike7, then demo.

Spike2 displays the file as it was last saved. The picture shows the state as shipped. You are looking at the raw data in the file. We call this a *time view*. It displays a time history of the data; the axis at the bottom is in seconds.



There are several *data channels* displayed in the window. Channel 1 holds a waveform, channel 2 holds events marked by dots, channel 3 holds events displayed as a mean frequency and channel 31 holds keyboard markers.

Selecting channels

Click a channel number to select a channel and highlight the channel number. Hold down the **Shift** key and click on a channel to select all channels between it and the last selection. Hold down **Ctrl** to select discontinuous channels. Some commands work on a list of selected channels (for example y axis optimisation). A red channel number means that the displayed data is modified, for example by a channel process or marker filter.

Toolbar and Status bar

The Toolbar is a shortcut to menu items. The Status bar provides information about your current activity. You can hide and show the bars with the **View** menu. You can also drag the Toolbar and “dock” it to any of the four sides. To find out what a Toolbar button does, leave the mouse pointer over the button for a few seconds.

There are other dockable toolbars. If you click the right-hand mouse button with the mouse pointer over an empty area of the Spike2 window or over a toolbar, a pop-up context menu appears from which you can show and hide any of the toolbars and the Status bar. Spike2 remembers the toolbar positions when the program closes and restores them on start up. Each user login has its own set of toolbar positions.



The bottom edge of the data window holds four buttons and a scroll bar. Try them.

The scroll bar controls movement through the file. You can also scroll the window one screen pixel with the **Left** and **Right** arrow keys. **Shift+Left** and **Shift+Right** move by several pixels and **Ctrl+Left** and **Ctrl+Right** move by half the screen width. The **Home** and **End** keys scroll to the start and end of the window.



Click this button or type **Ctrl+R** to halve the displayed time range (zoom in). The left hand edge of the display is fixed. **Ctrl+I** zooms in around the centre of the display.



Click this button or type **Ctrl+E** to double the displayed time range (zoom out). The left hand edge of the window is fixed unless the start plus the new width exceeds the file length, in which case the left edge moves back. If the new width would exceed the total length of the file, the entire file is displayed. **Ctrl+U** zooms around the centre.



Click this button to add a horizontal cursor to the display. You can delete and manipulate these cursors through the **Cursor** menu. Up to 9 horizontal cursors are allowed, numbered 1 to 9. These cursors are attached to the y axis of data channels. You can also add a particular cursor with the **Alt+1** to **Alt+9** keyboard shortcut.



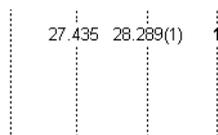
Click this button to add a vertical cursor to the display. Up to 10 vertical cursors can be present in a window. A cursor is a dashed line used to mark positions. You can remove cursors with the **Cursor** menu **Delete** command. You can add a cursor in these ways.

1. Click and release the button at the bottom left of the window (cursors 1 to 9).
2. Use the **Cursor** menu **New cursor** command (cursors 1 to 9).
3. Type **Ctrl+0** through **Ctrl+9** to show cursors 0 to 9 in the window.
4. Right click in the channel area and select **New Cursor** in the context menu.

You can scroll the screen to centre it on a cursor with the **Ctrl+Shift+0** through **Ctrl+Shift+9** key combinations. If the cursor does not exist, this has no effect.



Click the cursor button so that a vertical cursor is visible. Drag the cursor and observe how the mouse pointer changes. Use the **Cursor menu **Label mode** option.**



There are four labelling styles for the cursor: no label, position, position and cursor number, and number alone. You can select the most appropriate for your application with the **Cursor** menu **Label mode** option. To avoid confusion between the cursor number and position, Spike2 draws the cursor number in **bold** type when it appears alone, and in brackets when it appears with the position.

The mouse pointer changes when it is over a cursor into one of two possible shapes to indicate the two actions you can take with a cursor:



This shape indicates that you can drag the cursor from side to side. If you drag the cursor beyond the window edge, the window contents scroll to keep the cursor visible.



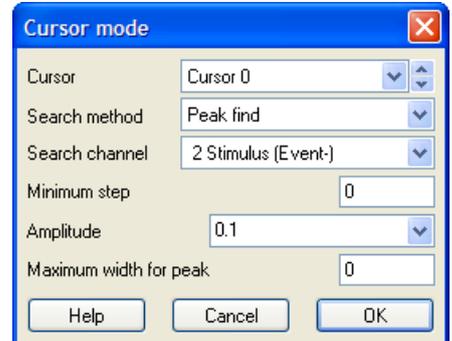
If you position the mouse pointer over the cursor label, the pointer changes to a 4-headed arrow to indicate that you can drag both the label and the cursor. This can be useful when preparing an image for publication and you need the cursor label to be clear of data. If you move the pointer to one side, or hold down the shift key, the pointer becomes a two-headed vertical arrow and you can drag the label, but not the cursor.

If you click the right mouse button with the mouse pointer over a vertical or horizontal cursor, the pop-up context menu includes commands to delete the cursor and also to set the cursor mode. Try deleting the cursors you have created. You can also delete one or all cursors from the **Cursor** menu.



Cursors can search for data features automatically. Use the Cursor menu Active modes command to open the Cursor mode dialog or right-click on a cursor and select Cursor mode from the context menu.

Vertical cursors can be Static or Active. A static cursor stays where you leave it. An active cursor can reposition itself by searching for user-defined data features. Cursor 0 is special; when active, you can command it to seek the next and previous data feature with keyboard commands. Cursors 1 to 9 are slaved to cursor 0; when cursor 0 moves, active cursors seek features in the order cursor 1 to 9.



Set cursor 0 to **Peak find** on channel 1. Set **Amplitude** to 0.1, **Minimum Step** to 0 and click **OK**.

Now try the `Ctrl+Shift+Right` key combination (`Right` is the right arrow key). Each time you press these keys, cursor 0 seeks the next peak on channel 1 that is at least 0.1 y axis units high. `Ctrl+Shift+Left` moves cursor 0 in the opposite direction.



Click on the data window then select one of the event channels (click the channel number) and zoom in so that only one or two events are visible on that channel. Type `Alt+Right` arrow.

If you have an event channel marking regions of interest, you can jump to the next or previous event by selecting the event channel, then using the `Alt+Right` arrow or `Alt+Left` arrow key combinations. Spike2 searches for the nearest event to the centre of the screen in the specified direction. If more than one event channel is selected, all selected channels are scanned for the nearest event.



Move the mouse pointer to the waveform channel, clear of any cursor. Click and drag a rectangle round a waveform feature, then release the button.

This action zooms the display so that the area within the rectangle expands to fill the entire waveform channel region. If your rectangle covers more than one channel, only the time axis expands. If your rectangle fits in one channel and has zero width, the y (vertical) axis changes to display the selected range and the time axis remains unchanged.



The mouse pointer changes to a magnifying glass when you hold the mouse button down in the data area to show that you are about to drag a rectangle or line to magnify the data.



If you hold down the control key before you hold the mouse button down, the mouse pointer changes to the un-magnify symbol. If you drag a rectangle, the rectangle holding the channel area shrinks to the rectangle you have dragged.

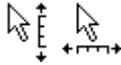
Whichever method you use to scale the data, you can return to the previous display with **Undo**. If you decide not to expand the display after starting to drag, return the mouse pointer to the original click position (so the rectangle has zero width and height). The rectangle will vanish and you can release the button without changing the display.



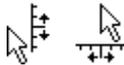
If you hold down the `Alt` key before you click and drag, Spike2 displays the size of the dragged rectangle next to the mouse pointer and does not zoom the display. If you release `Alt` but keep the mouse button down, you can use the `C` key to copy the current measurement to the clipboard and the `L` key to copy it to the Log view.



Move the mouse pointer over the x and y axes and experiment with clicking and dragging the axes. Try it with the **Ctrl** key held down.



When the cursor is over the tick marks of an axis, you can drag the axis. This maintains the current axis scaling and the axis moves to keep pace with the mouse pointer. You can do this with most x and y axes in Spike2. This is particularly useful for y axes as they do not have a vertical scroll bar. The window does not update until you release the mouse button. If you hold down the **Ctrl** key, the window will update continuously.



When the cursor is over the axis numbers, a click and drag changes the axis scaling. The effect depends on the position of zero on the axis. If the zero point is visible, the scaling is done around the zero point; the zero point is fixed and you drag the point you clicked towards or away from zero. If the zero point is not visible, the fixed point is the middle of the axis and you drag the point you clicked towards and away from the middle of the axis

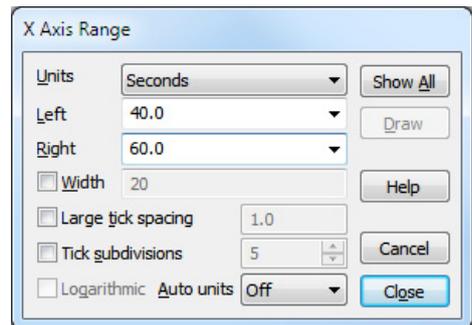
In a time view, result view, or XY view, you can drag the y axis so as to invert the axis. You are not allowed to invert the x axis. You are not allowed to invert or scroll the y axis in the spike shape dialogs.



Now double-click on the time (x) axis of the display to bring up the **X Axis Range** dialog box. Experiment with the settings to vary the time axis.

The Units field lets you draw the x axis in Seconds, milliseconds, hh:mm:ss (hours, minutes and seconds), or as Time of day.

The Left and Right fields set the window start and end times in the Units field units. The Width field shows the window width. Set the left and right positions, or check Width and set the Left and Width fields.



You can type new positions or use the drop down lists next to each field that give you access to cursor positions. The **Show All** button expands the time axis to display all the data. The **Draw** button updates the display to show the time range set by the Left, Right and Width fields.

You can enter times as `{{days:}hours:}minutes:}seconds` where `seconds` may include a decimal point and items in curly brackets are optional. Each colon promotes the number to the left of the colon from seconds to minutes to hours to days. If a file spans multiple days, the first day (regardless of the day of the month) is day 0, the second is day 1 and so on. With no colons, in `milliseconds` mode, values are in milliseconds. You can also enter a number in seconds, milliseconds or even microseconds by following it (without a space) by `s`, `ms` or `us`, for example `200ms`.

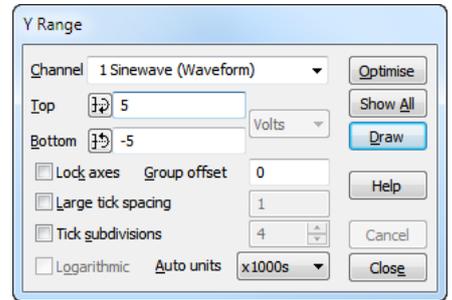
In addition to typing times, or selecting a time from the drop-down list, you can type in expressions using the maths symbols + (add), - (subtract), * (multiply) and / (divide). You can also use round brackets. For example, to display from 1 second before cursor 1 to one second past cursor 1 set **Left** to `Cursor(1)-1s` and **Right** to `Cursor(1)+1s`. **Draw** is disabled if any expression is invalid, or if the **Right** value is less than or equal to the **Left** value or if the new range is the same as the current range.

The **Large tick spacing** and **Tick subdivisions** fields let you customise the axis. Values that would produce an illegible axis are ignored. Changes to these fields cause the axis to change immediately; you do not need to click **Draw**. Click the **Help** button for more detailed information and for an explanation of the remaining fields.



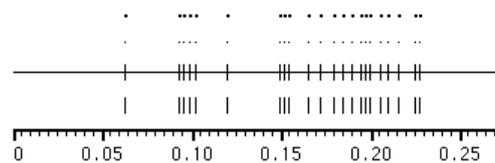
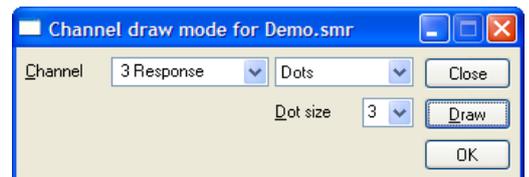
Now double-click on the y axis of the waveform channel to open the **Y Range** dialog. Experiment with adjusting the y axis ranges on different channels.

This dialog sets the y axis range of one or more channels. The **Channel** field is a drop down list; you can select a channel with a y axis, all channels with y axes, or all selected channels. You can **Optimise** the display, which makes sure that all the displayed data for the selected channel(s) fits in the y axis range, you can set the y axis limits as numbers, or you can use **Show All** to display the full range of the y axis. The **Lock axes** and **Group offset** fields are hidden unless you are working with grouped channels. You can also control the axis tick spacing, as for the x axis. Click **Help** for full dialog details.



Open the **View** menu **Channel Draw Mode** dialog. Experiment with different drawing modes for channel 3.

Spike2 files hold two basic data channel types: waveform and event. Waveform channels hold values representing the waveform amplitude at equal time intervals. Event channels hold the times at which something happened. There are four waveform drawing modes and many event drawing modes. For now, select channel 3 and try out the event modes.



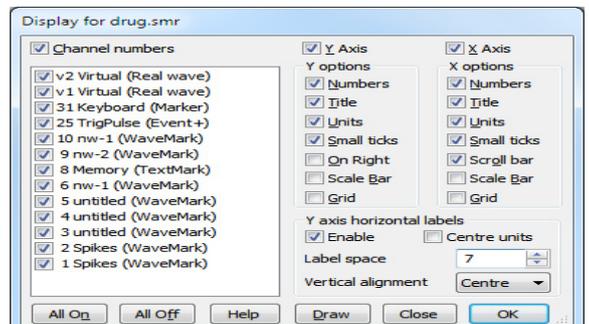
The simplest event mode is to draw the channel as dots. You can choose large or small dots (small dots can be difficult to see on some displays or when printed). You can also select **Lines** in place of **Dots**. The picture shows the result of

both types of display. If you select lines for a marker channel (like the keyboard channel), the marker codes are not drawn. Click the **Draw** button to cause an update without closing the dialog. See the *View menu* chapter for a description of all drawing modes.



Open the **View** menu **Show/Hide Channels** dialog. Experiment with the channels, axes and grid.

This dialog sets the channels to display in your window. Spike2 can handle up to 32 channels in a file, so the ability to display selected channels is quite important if you are to see any detail! The list on the left of the dialog holds all the channels that can be displayed.



You can show and hide the axes, grid and scroll bar in the window from this dialog and control the appearance of the x and y axes. Check the boxes next to the items for display and click the **Draw** button to see the result. The **Scale only** option draws axes as scale bars.



Make sure you have some cursors in the window. Use the **Cursor menu Display Y Values** command. Experiment with cursor positions and channel draw modes.

This is the Cursor Values dialog. The columns show the cursor times and data values for channels with a y axis. Channels with no axis display the time of the next event after the cursor. Check **Time zero** and **Y zero** for relative rather than absolute measurements.

Cursors	Cursor 0	Cursor 1	Cursor 2	Cursor 3
Time (s)	26.428747	30.807193	35.18564	39.564087
31 DupKey	27.03616	31.92832	37.04064	39.65696
1 Sinewave	2.5869173	-0.40142127	0.22451212	2.960858
3 Response	19.8036	18.549735	16.170691	16.120529
2 Stimulus	26.44303	30.82886	35.20149	39.61393
<input type="checkbox"/> Time Zero	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> Y Zero	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

You can choose which cursor to make the reference. The values at the reference cursor are unchanged, but the values at the other cursors have the value at the reference cursor subtracted. You can use this feature to show how data values have changed from a reference value.

If you move the cursors or change the channel display mode, the values in this window update to reflect the change of position. Likewise, if you add or remove data channels or cursors from the display, the cursor dialog will change.

You can select fields in this dialog and copy them to the clipboard. Click on a field to select it or drag across the data area to make a rectangular selection of fields. Click at the top or left hand edge to select an entire column or row. Click in the top left hand box to select all the fields. Hold down the control key and click at the top or left hand edge for non-contiguous selection of rows or columns.



Now open the **Cursor menu Cursor Regions** dialog. Experiment with changing cursor positions and channel display types.

The regions dialog looks at the data values between cursors. There are many modes of operation, for example **Area**, **Mean** and **Slope**. See the *Cursor menu* chapter for a full list. To change the mode, click in the box at the bottom left of the dialog (holding “Area” in the picture) and you will see a drop down list of all available modes. You can also make measurements relative to one of the regions by checking the **Zero region** box and choosing a reference region.

Cursors	0 - 1	1 - 2	2 - 3
Time (s)	5.0247856	4.3784467	4.3784467
1/Time (Hz)	0.19901347	0.2283915	0.2283915
31 DupKey	0.39802738	0.45678265	0.22839133
1 Sinewave	1.3717608	1.3835991	1.3982419
3 Response	17.862161	21.821886	14.582918
2 Stimulus	13.532931	18.728089	18.271306
<input type="checkbox"/> Zero Region	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Mean			

For a waveform channel, the **Area** is the area between the waveform and the y zero level, the **Mean** is the mean level of the signal. If you select **Slope**, Spike2 calculates and displays the gradient of the least-squares best-fit line to the data.

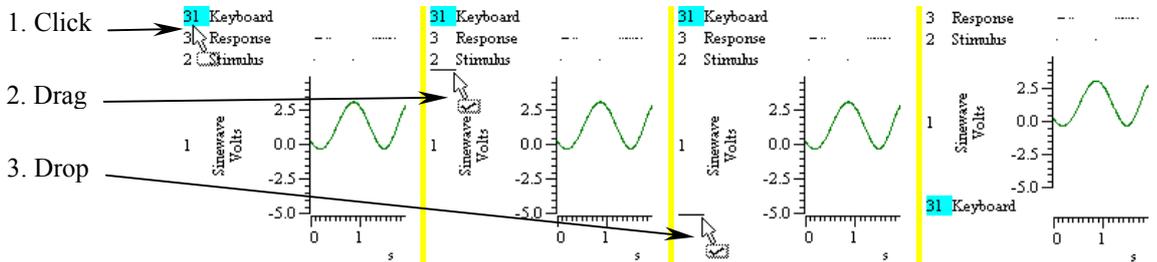
For an event channel, the **Area** is the number of events in the region. **Mean** is the count of events divided by the width of the region. **Slope** has no meaning for an event channel.

Long drawing and calculation times

Spike2 can display huge data files (sizes of hundreds of megabytes and more). However, it takes time to move such large quantities of data from disk into memory. If you accidentally try to draw a huge file, or use the **Cursor regions** dialog to calculate a value from a huge file, you may end up waiting for a long time. You can break out of such operations with the **Ctrl+Break** key combination.



Use the View menu Standard Display command. Click on the Keyboard channel number (31) and drag it down over the other channel numbers.



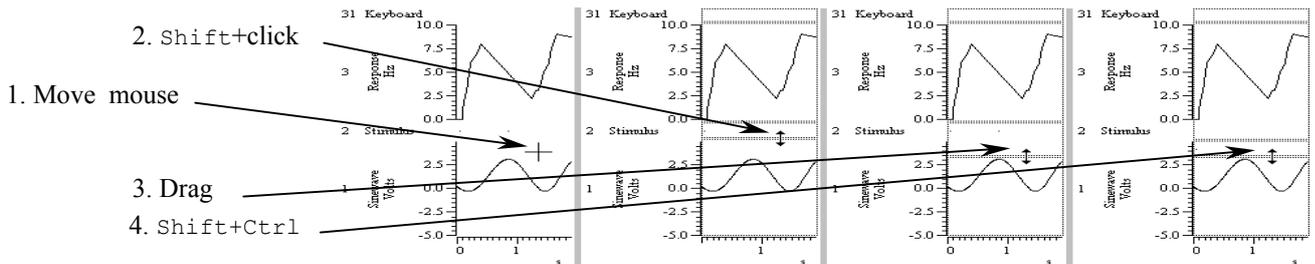
As the mouse pointer passes over each channel, a horizontal line appears above or below the channel. This horizontal line shows where the selected channel will be dropped. Drag until you have a horizontal line below channel 1 and release the mouse button. Channel 31 will now move to the bottom of the channel list. Type `Ctrl+Z` or use the Edit menu Undo to remove your change.

You can move more than one channel at a time. Spike2 moves all the channels that are selected when you start the drag operation. For example, hold down `Ctrl` and click on the channel 3 number. Keep `Ctrl` down and click and drag the channel 2 number. When you release, both channels will move. The mouse pointer shows a tick when you are in a position where dropping will work.

The usual Spike2 channel order is with low numbers at the bottom of the screen. If you prefer low numbers at the top of the screen, open the Edit menu Preferences and check Standard Display shows lowest numbered channel at the top, then use the View menu Standard Display command.

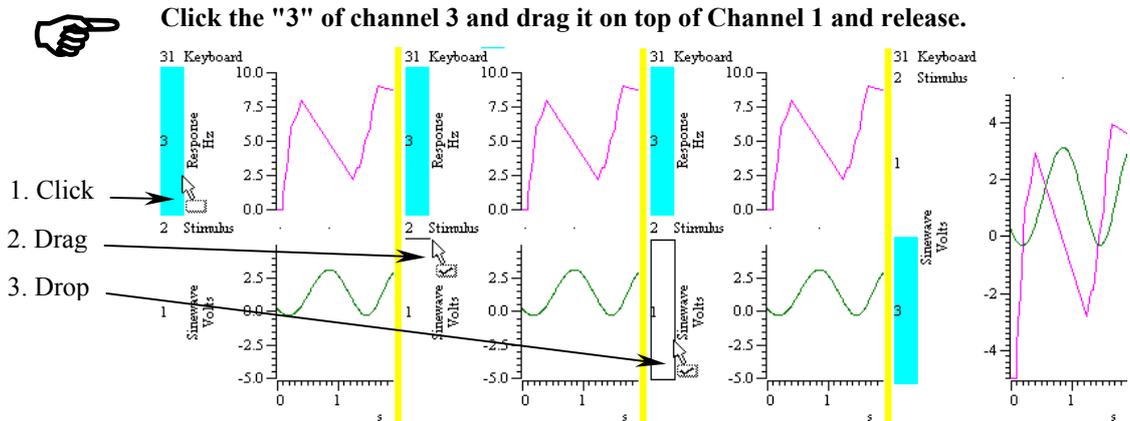


Change the channel 3 drawing mode to Mean frequency. Hold down the Shift key and move the mouse over the data area. Hold the Shift key down and click. Drag up and down and release the mouse.



When you click with `Shift` down, the mouse jumps to the nearest channel boundary and you can change the boundary position by dragging. With `Shift` down, you can move the edge up and down as far as the next channel edge. You can undo changes or use Standard Display to restore normal sizes.

If you add `Ctrl`, all channels with a Y axis are scaled. If there are no channels with a Y axis above or below the drag point, then all channels scale. You can force all channels to scale by lifting your finger off the `Shift` key (leaving `Ctrl` down) after you start to drag the boundary.



The channels now share the same space with the channel numbers stacked up next to the y axis. The visible y axis is for the top channel number in the stack. To move a stacked channel to the top, double-click the channel number. To remove a channel, drag the channel number to a new position.

When you drag channels, and at least one of the selected channels has a Y axis, you can drop the channels with a Y axis on top of another channel with a Y axis. As you drag, a hollow rectangle appears around suitable dropping zones. You can also drop between channels when a horizontal line appears.

The channel that owns the Y axis is drawn last. Put channels that fill in areas, such as sonograms or events drawn in rate mode, at the bottom of the stack, as they mask channels below them. Stacked channels keep their own y axes and scaling. You can force them to use the displayed y axis range by checking the Lock box in the Y Range dialog.

Result views

So far, you have been looking at windows holding raw, unprocessed data. We call these *Time views*. There is another type of data window, called a *Result view*, which holds the result of analysing time view data. There are two steps in the analysis:

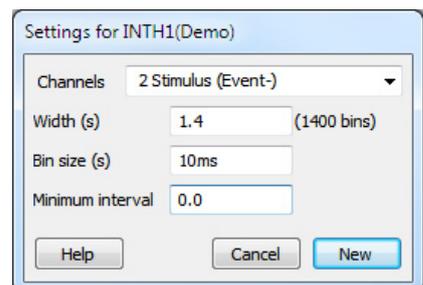
1. Set the type of analysis, the channels to use, the width of the analysis or bins to generate and any other parameters. This creates a new, empty result window.
2. Set a time range in the time view. Spike2 processes the data and adds it to the result.

Repeat step two as often as is required to accumulate results from different sections of data. Most analyses generate windows that can be thought of as histograms or waveforms, that is a list (array) of data values equally spaced in the x direction.



Make the original time view of the data the current window by clicking on it. You may find it easier if you close all the other windows first. Use the Analysis menu New Result View command to select an Interval Histogram.

The dialog is prompting you for information to define a new window. There are four fields to fill in that define the interval histogram. The first field selects channels to analyse. You can select any channel that does not hold waveform data; the drop down list only includes suitable channels. If you prefer, you can type in a channel list made up from channel numbers and channel ranges separated by commas. A channel range is a list of consecutive channels written as



the first and last channel separated by two dots. For example, 1..3 is the same as 1,2,3 and 3..1 is the same as 3,2,1. You can combine the two, 1,3,4..8 is a valid channel list.

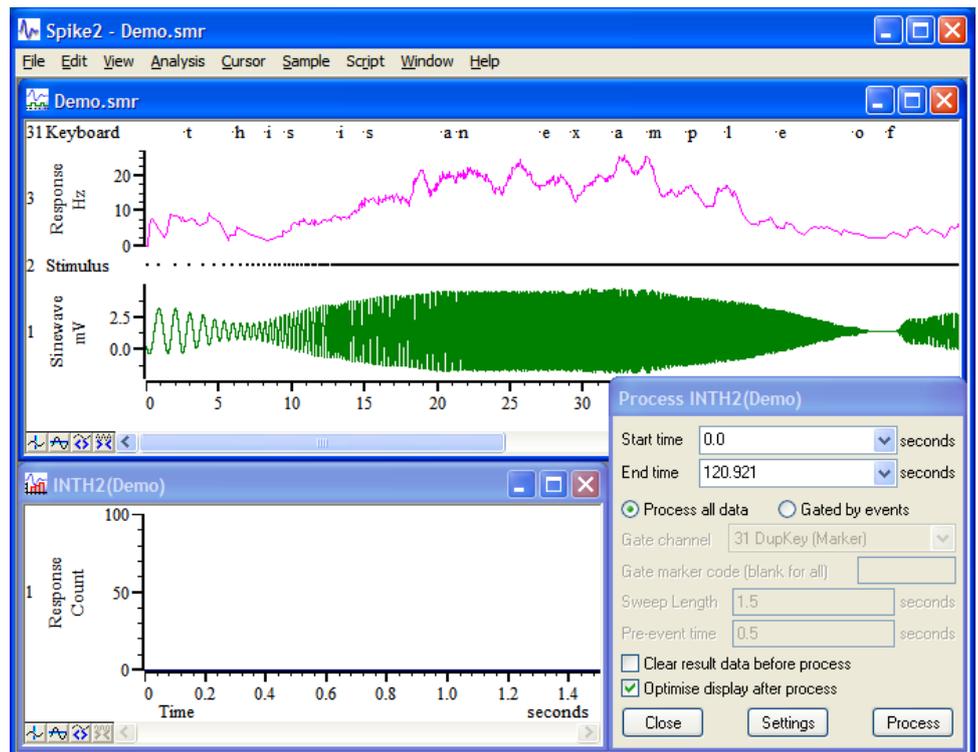
The next two fields set the histogram width and bin size in seconds. Spike2 calculates the number of bins and displays them to the right of Width. The width is limited only by the available memory to store the bins. The histogram bins must be at least one Spike2 clock tick wide (you set the clock tick size when configuring the data sampling and it is usually in the range 2 to 50 μ s). The bin width is rounded to a multiple of this clock tick, so may not stay exactly as you set it. The last field sets the minimum interval that appears in the new window. In most cases you will leave this set to 0.0 seconds.

If you prefer, you can override the seconds units (as is done for the Bin size field in the example) and type a number followed immediately by ms or us to force the use of milliseconds or microseconds.

The first thing to do is to select a channel to analyse. A suitable one for our purposes is channel 2, so select this one now using the drop down list. Set the remaining fields as they are in the picture above; 1.5 seconds wide, 0.01 seconds bin width (10 milliseconds), and a minimum interval of 0.



Once you have set these values, click the **New** button to generate the new result window. Now set the region of data to analyse.



When you click the **New** button Spike2 creates a new window ready to display the result of the analysis, the setup dialog vanishes and the **Process** dialog appears. You must now set the region of the data document to analyse with the **Start time** and **End time** fields.

You can choose to process all the data in the time range, or you can process only data that lies within a specified time range of event markers on a specified channel. Normally you will want to process all the data, so you should select **Process all data**.

The two check boxes in the **Process** dialog determine how to treat the result of the analysis. You can choose to clear the result window contents before you analyse the data, otherwise each new result is added to the previous one. You can also choose to optimise the result view display after each analysis so that the full range of the data is visible.

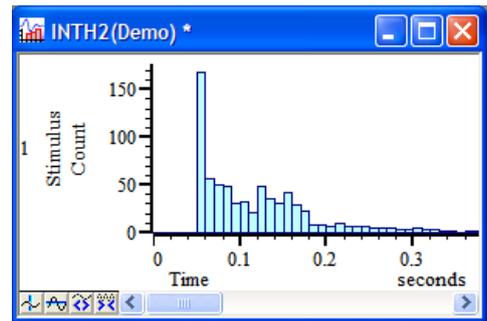
If you decide that you have not set the original parameters correctly, you can click on the **Settings** button to go back to the previous step and correct the values.



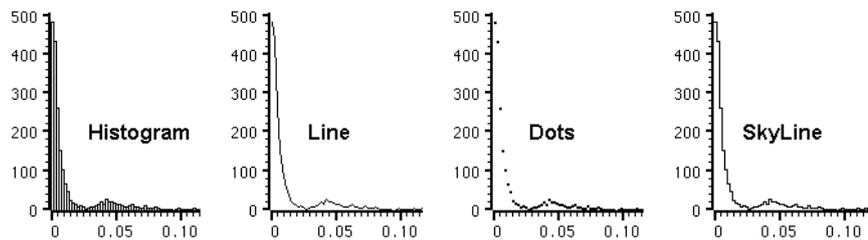
When you have set the time region, click the **Process button.**

The dialog vanishes and the result window shows the analysed data. The picture shows the data with the zoom-in button clicked twice.

You can recall the analysis dialog by selecting the **Process** command from the **Analysis** menu. Do this now and click the **Process** button again. The data in the result window will double in size (as long as you have not checked the **Clear result view before process** box).



Experiment with the **Channel Draw Mode command in the **View** menu.**



There are five drawing styles available for result window histograms: **Histogram**, **Line**, **Dots**, **SkyLine** and **Cubic Spline**. These styles are self-explanatory. The various analysis routines that create result windows will select an appropriate style.



Experiment with cursors in this new window.

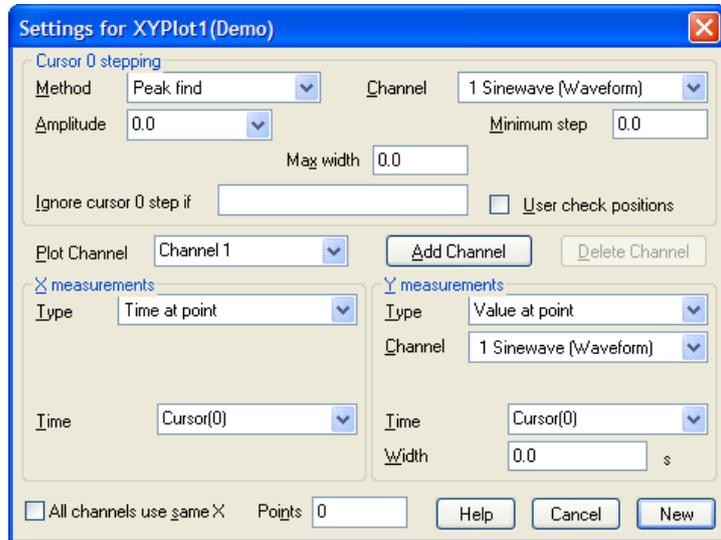
You will find that the cursors behave in a very similar manner to the original time window. You can create and position up to 9 vertical cursors and 4 horizontal cursors and display the **Cursor regions** and **Cursor values** dialogs from the **Cursor** menu. However, there are no active cursors in a result window and no vertical cursor 0.

XY views

In addition to **Time** and **result** views, there are **XY** views. These hold up to 256 data channels that share the same x and y axes. Each channel is a list of (x,y) co-ordinates and has its own point marking style, line style and colour. **XY** views are commonly used to draw trend plots of data extracted from time views and for user-defined images and graphs generated by a script. We will start by generating a trend plot.

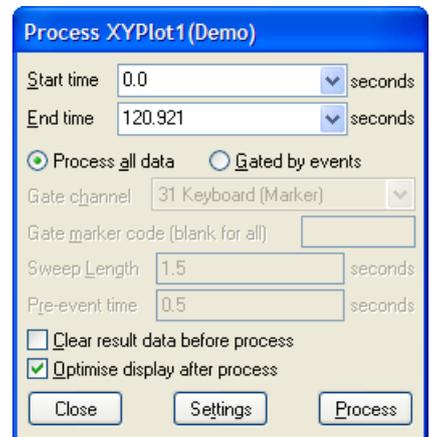


Close all windows except demo.smr. Set cursor 0 to **Peak find** mode as you did earlier, then use the **Analysis menu Measurements->XY view** command.



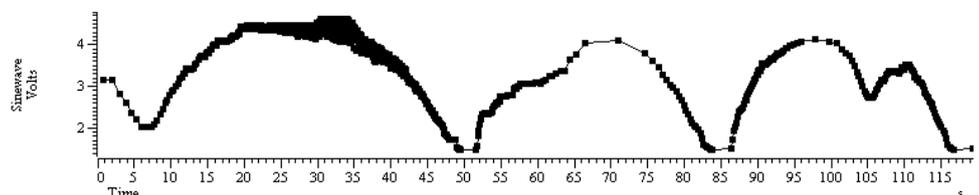
In a trend plot, cursor 0 steps through a data region. At each step, x and y values are derived from the data and added to an XY view. We will plot the peak heights of the waveform channel in the demo.smr file. If you set up cursor 0 correctly, the Cursor 0 stepping section of your dialog will match the picture; adjust it to match if it differs.

Set the X measurement area and the Y measurement area to match the picture. Set the Type fields first, then adjust the remaining fields and finally click the New button. A new XY window opens together with the Process dialog. The Start time and End time fields set the limits for cursor 0 stepping.



In the same way as for the interval histogram, you can either process all data in the time range or you can choose to process only data that is associated with particular events. You will usually want to select Process all data.

In this case we want to step through the entire file. Now click the Process button.

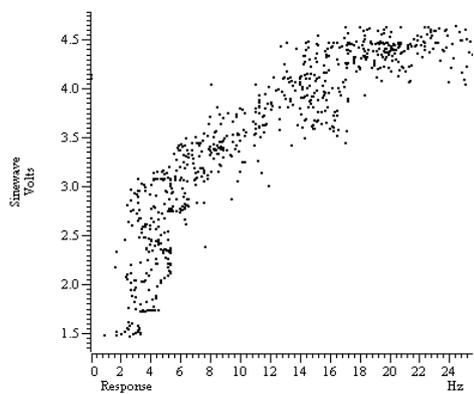
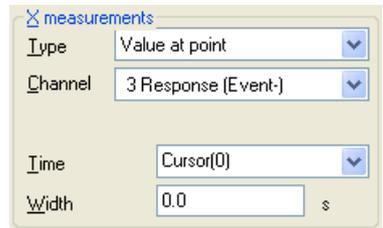


The result looks a little blotchy! Double click the data area of the XY window to open the XY Drawing mode dialog. Set the Size field to 0 and click Apply to hide the markers at each data point. You can use a variety of drawing styles. When you have finished, click OK.



So, what have we done? In each case our x measurement was set to a time, and that time was the cursor 0 position. The y measurement was the value of channel 1 at a time, again being the cursor 0 position. Cursor 0 was set to step through the data on channel 1 looking for peaks in the data. The result is a graph of peaks of each cycle of data.

Suppose we suspected that the event rate of channel 3 is related to the amplitude of the peaks. To test this we will change our plot. Click on the `demo.smr` window and change the drawing mode of channel 3 to Mean frequency over 1 second. Right click on the XY view and select **Process settings**. Change the Type to Value at point, select channel 3, then click the Change button. You will be warned that this will delete the previous result, so say this is OK.



The result will look a little messy with lines joining the points, so activate the XY Drawing mode dialog and set the Join style to Not joined and the marker size to 2. The result should look something like the picture.

You will find detailed information about all the different measurements you can make in a trend plot in the *Analysis menu* chapter. The *Cursor menu* chapter has details of the active cursors.



Use the Script menu Run Script command and select the Load and run... command. Locate the `scripts` folder (in the folder where you installed Spike2), and open the file `clock.s2s`.

Spike2 will load and run this script. It generates an analogue clock in an XY view. You can move and resize the clock window. You can stop the script running (and regain control of Spike2) by clicking on the OK button at the upper right hand side of the Spike2 window. You can read more about XY views in the script language manual and in the *Spike2 Training Course* manual.

Summary

If you have followed this chapter, you will now be familiar with the basic actions required to use Spike2. The next chapter contains useful general information about Spike2 and the following one deals with the special actions required to configure the system for sampling your own data. The remainder of the manual covers the menu commands, the memory buffer, the sequenced output system, cut and paste to other applications, printing, the spike shape capture and editing module, spike clustering, digital filtering and utility programs for fixing damaged files and testing your 1401.

If you have a Spike2 supported programmable signal conditioner, see the *Programmable Signal Conditioners* chapter. Standard 1401 or 1401plus users with the optional 1401-18 discriminator card should see the *1401-18 Programmable Discriminator* chapter.

The *Script menu* chapter describes the menu options that control the script system. The script language itself is not covered in this manual; see the companion text *The Spike2 script language* for a full description.

General information

This chapter gathers together information that would otherwise be scattered and repeated throughout the manual. Channel lists are used in many dialogs and also in the script language. Expressions can be used in many dialogs where x axis values are wanted and also more generally. There are many keyboard shortcuts in Spike2; they are gathered together here. The command line lets you control Spike2 from other programs; script users can even launch another copy of Spike2.

Channel lists

In many places where you are prompted for a channel, you can select a channel from a drop down list, or you can type a channel list. A channel list is a list of channel numbers or channel ranges separated by commas. A channel range is two channel numbers separated by two periods or a hyphen, for example 4..7, which is equivalent to channels 4, 5, 6 and 7. The channel range 7-4 is equivalent to channels 7, 6, 5 and 4. The following channel list:

```
1,3..5,7
```

means channels 1, 3, 4, 5 and 7. In most cases, Spike2 checks channel lists and removes channels that are not suitable for the operation. For example, if you open the `demo.smr` file supplied with Spike2, select an Interval histogram and type a channel list of 1..32 and then click on another field, Spike2 reformats the list to 2,3,31 as these are the only suitable channels. It is not an error for a channel list to include unsuitable channels, however it is an error for a channel list to include no suitable channels.

You can replace channel numbers with symbolic names, for example `m1` for a memory channel or `v2a` for the first duplicate of virtual channel 2. If you specify a range of duplicates of the same channel, the range expands by duplicates, not by channel numbers. For example: `2..2c` means 2, 2a, 2b, 2c. We allow up to 52 duplicates of a channel using upper case `A..Z` for the second set of 26. So `2..2C` is the same as `2..2z, 2A..2C`.

Channel lists can also be used in script commands, for example: `ChanShow("1..4")`. Script commands that will accept this format describe the argument as `cSpC`. You can find more information about channel specifications in the script language documentation.

Dialog expressions

Many dialogs in Spike2 accept an expression in place of a number. These expressions can be divided into two types: *numeric expressions* and *view-based expressions*.

Numeric expressions

A numeric expression is composed of numbers, the arithmetic operators `+`, `-`, `*` and `/`, the logical operators `<`, `<=`, `=`, `>=`, `<>` and `?` and round brackets `(` and `)`. The result of a logical comparison is 1 if the result is true and 0 if the result is false. You may not have come across `?` which is used as:

```
expr1 ? expr2 : expr3
```

The symbols `expr1`, `expr2` and `expr3` stand for numerical expressions. The result is `expr2` if `expr1` evaluates to a non-zero value and `expr3` if `expr1` evaluates to zero. This can be used in the Cursor mode dialog to give a cursor position if a search fails based on some other information, for example:

```
Cursor(2)>Cursor(1) ? Cursor(2) : Cursor(1)
```

This evaluates to the position of the rightmost of cursors 1 and 2.

If you write expressions involving more than one operator, for example `1+2*3` you need to know if this is evaluated as `(1+2)*3` or as `1+(2*3)`. This is determined by the operator precedence level (described below).

View-based expressions These expressions follows the rules for numeric expressions and allow references to positions along the x axis. If a dialog field is documented as allowing expressions, and the field supplies an x axis position (for example a time), then you can use the following:

<code>Cursor(n)</code>	Where n is 0 to 9 returns the position of the cursor. If the cursor does not exist or the position is invalid, the expression evaluation fails.
<code>C0 to C9</code>	This is shorthand for <code>Cursor(0)</code> to <code>Cursor(9)</code> .
<code>XLow()</code>	The left hand end of the visible x axis in seconds for a time view, bins for a result view, and x axis units for a XY view.
<code>XHigh()</code>	The right hand end of the visible x axis.
<code>MaxTime()</code>	The right hand end of the time axis in seconds for a time view and bins for a result view. It is not valid in an XY view.
<code>MaxTime(n)</code>	The time of the last data item on channel n in a time view.

You can add `View(-1)` . before these expressions to force the expression to be evaluated for the time view linked to the current view, for example `View(-1).Cursor(0)`. This is required when the current view is a result view and you wish to access timing information from the time view that the result view is based on.

Times as numbers Internally, all times are in seconds. However, where a time is typed into a dialog you can use `{{{days:}hours:}minutes:}seconds` where `seconds` may include a decimal point and items enclosed in curly brackets are optional. Each colon promotes the number to the left of the colon from seconds to minutes to hours to days. Times may only contain numbers and colons, white space is not allowed. Alternatively, a number with no colons can have a `s`, `ms` or `us` suffix to indicate: seconds, milliseconds or microseconds. So the following are equivalent: `1.6`, `1.6s`, `00:00:01.6`, `1600ms`, `1600000us`. The X Axis dialog allows entry of times in milliseconds if the **Units** field is set to **milliseconds**.

Operator precedence In the table, **LHS** means the value of the expression to the left of the operator as far as the next operator of same or lower precedence, **RHS** means the value of the expression on the right up to the next operator of the same or lower precedence. Where operators have the same level, evaluation is from left to right. The order from high to low is:

Level	Name	Return value
5	() Brackets	Everything inside a pair of brackets is evaluated before considering the effect of an adjacent operator.
4	* Multiply / Divide	LHS multiplied by RHS LHS divided by RHS. It is an error for RHS to be zero
3	+ Add - Subtract	LHS plus RHS LHS minus RHS
2	< Less than <= Less or equal = Equal >= Greater or equal > Greater than	If LHS less than RHS then 1 else 0 If LHS less than or equal to RHS then 1 else 0 If LHS equal to RHS then 1 else 0 If LHS greater than or equal to RHS then 1 else 0 If LHS greater than RHS then 1 else 0
1	? Ternary operator	LHS?A : B has the value A if LHS is not 0, and B if it is 0. Put spaces around the colon to distinguish it from a time.

`1+2*3` has the value 7 because multiply has a higher precedence level than add.

Script language compatibility The expressions are compatible with the script language except for use of `C0` to `C9` and `H1` to `H9` as shorthand for `Cursor(0)` to `Cursor(9)` and `HCursor(1)` to `HCursor(9)` and the use of colons, `s`, `ms` and `us` to denote times. If you use these in a script you will get syntax errors. However, you can use these constructs in strings passed as expressions to `CursorActive()` or `MeasureChan()`.

Data view keyboard shortcuts

The following shortcut key combinations can be used in a time, result view or XY view except for `Ctrl+Shift+Left/Right` and `Alt+Left/Right`, which are only available in a time view.

Key	Operation
<code>Left/Right</code>	The two arrow keys scroll 1 pixel left or right.
<code>Shift+Left</code>	Scroll several pixels left.
<code>Shift+Right</code>	Scroll several pixels right.
<code>Ctrl+Left</code>	Scroll half a screen left.
<code>Ctrl+Right</code>	Scroll half a screen right.
<code>Ctrl+Shift+Left/Right</code>	Time views only. If cursor 0 is active, search for the next/previous feature and scroll the screen to make it visible.
<code>Alt+Left</code> <code>Alt+Right</code>	Time views only. Search selected event channels for the previous/next event that is nearest to the screen centre and make it the screen centre or make a sound if there are no selected channels or no more events.
<code>Alt+Shift+Left/Right</code>	Time views only. Jump to the next display trigger point (only if the display trigger is enabled).
<code>Home/End</code>	Scroll to the start/end of the data.
<code>Ctrl+n</code>	Where <i>n</i> is 0 to 9. Fetch vertical cursor 0 to 9. If the cursor does not exist it is created. Cursor 0 exists only in time views.
<code>Alt+n</code>	Fetch or create horizontal cursor <i>n</i> (1 to 9).
<code>Ctrl+Shift+n</code>	Where <i>n</i> is 0 to 9. Centre the window on the cursor, if it exists.
<code>Ctrl+A</code>	Select all channels. If all channels are selected, unselect all channels.
<code>Ctrl+C</code>	Copy the image of the view to the clipboard.
<code>Ctrl+E</code>	Expand (zoom out) the data view around the left edge of the window.
<code>Ctrl+I</code>	Reduce (zoom In) the data view around the centre of the window.
<code>Ctrl+K</code>	Show the x axis dialog.
<code>Ctrl+N</code>	Open a new data file.
<code>Ctrl+O</code>	Open the file open dialog.
<code>Ctrl+P</code>	Print the current data file.
<code>Ctrl+Q</code>	Optimise selected channels. If none selected, optimise all channels.
<code>Ctrl+L</code>	Open the Evaluate window to run single line script commands.
<code>Ctrl+R</code>	Reduce (zoom in) the data view around the left edge of the window.
<code>Ctrl+T</code>	Create a TextMark during sampling (if the TextMark channel exists).
<code>Ctrl+U</code>	Expand (zoom out/Up) the data view around the centre of the window.
<code>Ctrl+Y</code>	Show the y axis dialog.
<code>Ctrl+Z</code>	Undo the last undoable operation.
<code>Ctrl+Break</code>	Break out of long drawing or calculation operations.

Text view keyboard shortcuts

Text views have more keyboard short cuts than any other area of Spike2. We have grouped them by function to make the huge list more digestible.

Text caret control

The text caret is a flashing vertical bar that indicates the current position. Do not confuse this with the I-beam mouse pointer which does not flash and which indicates the mouse position. Each time you click and release the left mouse button (we assume you haven't swapped the mouse buttons), the caret moves to the nearest character position to the click point. To select text with the mouse, click at one end of the text you want to select and

drag (move the mouse with the button held down) to the other end of the text. You can also use the keyboard to move the caret and select text:

Key	Operation (+Shift to extend a text selection)
Left arrow	Move the caret one character to the left. At the start of a line it wraps to the end of the previous line.
Right arrow	Move the text caret one character right. You can move it into uncharted territory beyond the end of the line. It does not wrap to the next line.
Up arrow	Move up one line.
Down arrow	Move down one line.
Ctrl+Left Ctrl+Right	Move one word to the left/right. Words are defined to be useful when operating on scripts.
End	Move the caret to the right of the last character on the line.
Home	Move the text caret to the start of the current line.
Ctrl+End	Move the text caret to the right of the last character in the file.
Ctrl+Home	Move the text caret to the left of the first character in the file.
Ctrl+] ([)	Start of next (previous) paragraph (after empty line)
Ctrl+\ (/)	Word part right (left).
Insert	Swap between insert mode caret and overwrite caret _

Cut, Copy, Paste, Delete, Undo and Redo

Some of these operations are also available from the Edit menu and the main toolbar.

Key	Operation
Ctrl+A	Select all the text in the document.
Ctrl+C Ctrl+Insert	Copy selected text to the clipboard. If no text is selected, nothing is copied. Some keyboards have <i>Ins</i> in place of <i>Insert</i> .
Ctrl+Shift+T	Copy the current line to the clipboard.
Ctrl+V Shift+Insert	Paste the contents of the clipboard into the text at the caret. If there is a selection, the selection is replaced.
Ctrl+D	Duplicate the selection.
Ctrl+X	Cut the selected text and copy it to the clipboard.
Shift+Del	Cut the selected text and copy it to the clipboard.
Backspace	Delete the selection or the character to the left of the text caret.
Del	Delete the selection or the character to the right of the text caret.
Ctrl+Del	Delete word right. Add <i>Shift</i> to delete to the end of the line.
Ctrl+D	Duplicate the selection.
Ctrl+Shift+L	Delete the current line.
Ctrl+Z, Alt+ Backspace	Undo the last interactive text operation. The editor supports more or less unlimited levels of Undo.
Shift+Ctrl+Z	Redo the immediately previous Undo operation.

Miscellaneous

These commands do not fit into any other category!

Key	Operation
Ctrl+U	Convert the selection to upper case. Add <i>Shift</i> for lower case
Ctrl+Add, Sub	Zoom text (change font size). Add and Sub are numeric keypad + and - keys. You can also zoom the text using <i>Ctrl</i> +the mouse wheel.

Find, Replace and Bookmarks

The Find and Replace commands can be accessed from the Edit menu, from the Edit Toolbar and by keyboard short cuts:

Key	Operation
Ctrl+F	Open the Edit menu Find dialog. In addition to searching for text you can also use this dialog to bookmark all matching text.
Ctrl+H	This shortcut key opens the Edit menu Replace dialog.
F3	Repeat the last find operation in the same direction. You can use the toolbar to search forwards or backwards.
F2	Move the text caret to the next bookmark. You can use the edit toolbar to move to the next or previous bookmark.
Ctrl+F2	Toggle bookmark on the current line. You can use the edit toolbar to set or clear a bookmark and to clear all bookmarks.

Bookmarks tag a line for future reference. They are displayed as a blue mark to the left of the text. Bookmarks are kept as long as the current file is open; they are lost when you close the file. The easiest way to use a bookmark is from the Edit Toolbar. You can show and hide this from the Edit menu (when a text-based window is active), or by clicking the right mouse button on any toolbar or on the Spike2 application title bar and using the pop-up context menu that appears.

Indent and Outdent

The structure of Spike2 scripts is often made clearer by indenting program structures. To make this easier, you can indent and outdent selected blocks of text to the next or previous tab stops. The tab size is set in the Edit menu Preferences option.

Key	Operation
Tab	If there is a multi-line selection, all lines included in the selection are indented so that the first non-white space character is at the next tab stop. If there is no selection, a tab character is inserted (or spaces to the next tab stop depending on the Edit menu Preferences settings).
Shift+Tab	If there is a multi-line selection, all selected lines are out-dented so that the first non-white space character on the line is at the previous tab stop. If there is no selection, the text caret moves to the previous tab stop unless it is already at one.

Drag and drop

The editor supports drag and drop of text both within Spike2 and between Spike2 and other applications that support it (for example the Spike2 Help system). Spike2 also supports drag and drop for rectangular text areas.

Operation	Method
Move block	Select the text to move. Move the mouse pointer over the selected text and hold down the left mouse button and drag. The mouse pointer will indicate that you can now drag the text and the text caret will show the insertion point. Drag the text to the desired insertion point and release.
Copy block	Select the text to copy. Hold down the Ctrl key and move the mouse pointer over the selected text, click and drag. A small + symbol indicates the copy operation and the text caret will show the insertion point for the duplicate. Drag the text to the target position and release the mouse button to duplicate the text. The Ctrl key must be down when you release the mouse button or the operation will move the text.

Select rectangular text area You can select, cut, paste and drag rectangular selections within Spike2. To select a rectangular area hold down the `Alt` key then select text with the mouse. The point where you hold down the mouse button will be one corner of the selection, the point where you release the mouse will be the other corner. You can use this feature to change the alignment of comments in a script, or to convert a single column of numbers into multiple columns. You can also paste such text into other applications as plain text.

The Spike2 command line When Spike2 starts, it checks the command line for option switches and for files to load. If there is no command line, Spike2 looks in the folder that Spike2 ran from for a script called `startup.s2s` and runs it if it exists. If `startup.s2s` is used or the command line loads a file, any start up messages that wait for a user response are suppressed.

The command line holds options and file names separated by white space characters (space and tab). If a file name contains spaces, you must surround the file name with quotation marks. Options start with `/` or `-` followed by a character to identify the option.

`/M` When you start Spike2, it checks if there is already a running copy. If there is, the new one quits. This option removes the check, allowing multiple copies to run on a single system. You need a Spike2 licence for each copy except when using multiple synchronised 1401s to capture related data on one computer under the control of a single operator, when one licence is sufficient. To do this, you must set separate file names for each 1401 in the Automation tab of the Sampling Configuration dialog.

`/Un n` is 1-8 to select a 1401 when you have more than 1. The default is `/U0`, which uses the lowest-numbered unused 1401. You set a device number in the CED 1401 device settings in the Device Manager (My Computer->Properties->Hardware).

`/Q` Quiet startup. Suppress all message boxes and the Spike2 "splash screen".

The remaining items in the command line are assumed to be file names. Spike2 attempts to load the files in command line order (from left to right). The files must have extensions so that the file type is known. If a script file is included in the command line, Spike2 runs it before continuing with the remainder of the command line.

As an example, suppose we want to launch Spike2 so that it automatically opens a data file called `lots of data.smr` and runs `doit.s2s` to process it. Follow these steps:

1. Create a short cut to `sonview.exe` (this is the Spike2 program).
2. Right-click on the new short cut and select **Properties** and open the **Shortcut** tab.
3. Add "`lots of data.smr`" `doit.s2s` to the end of the **Target** field.
4. Set the **Start in** field to the folder that contains your files.
5. Click **OK**.

This example assumes that both files are in the same folder. You could also have included the full path to each file in the command line.

Shell extensions The Spike2 installation adds shell extensions that display additional information in Windows Explorer when the mouse pointer hovers over Spike2 data and script files. In Windows NT2000 and XP you can also display file comments. To do this, set **Details** display mode, then right click in the column headers and check **Comments**.

These shell extensions are automatically removed if you uninstall Spike2. To remove the shell extensions manually, open a command prompt window and type:

```
C:\>cd \Spike7 Change to the Spike2 folder
C:\Spike7>regsvr32 /u SonInfo For all 32-bit Windows versions
C:\Spike7>regsvr32 /u SonCols Only for XP (32-bit) or later
```

Sampling data

Sampling introduction

If you worked through the *Getting started* section you already have most of the skills to sample a new data document. Sampling a new document is the same as working with an old document, except that the new document grows in length.

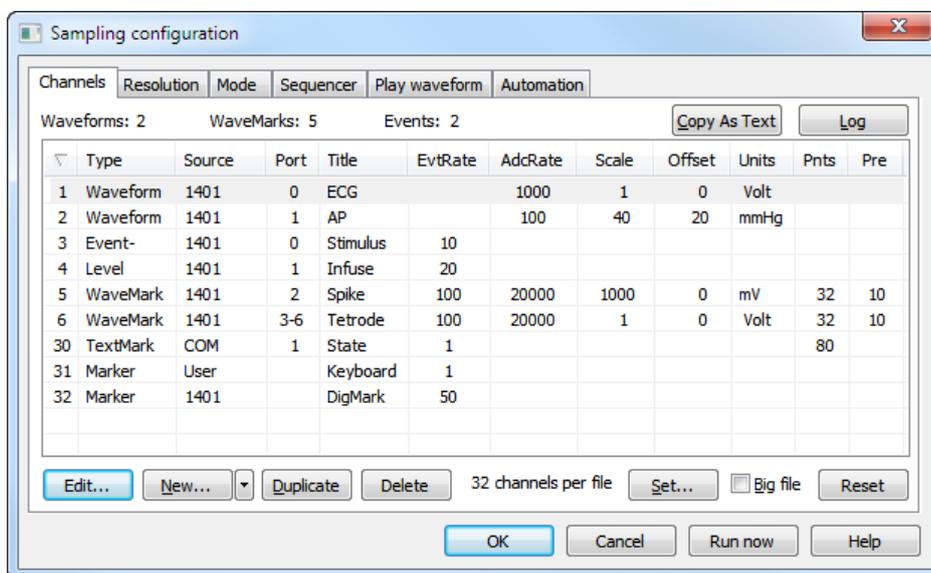
Sampling configuration



Before starting to sample data you must set the sampling configuration in the **Sample** menu **Sampling Configuration** dialog. This is a tabbed dialog; it holds several pages activated by the tabs at the top. Each page controls a different aspect of data capture. The title bar holds the name of the configuration file from which the configuration was read. There is a * at the end of the name if the configuration has been changed. The **OK** button accepts the current state, **Cancel** rejects changes made since you opened the dialog and **Run now** opens a new data document, ready to sample. You can also open this dialog from the toolbar. You can resize the dialog by dragging the top or bottom edge.

Channels

The **Channels** tab lists the channels to sample. One channel is always selected. The left-hand column is the channel number; the remaining columns list the channel type, the data source and physical port, the title, then data that varies with the channel type. You can reorder the items by clicking on the headers of most of the columns.



The top line displays the count of Waveform channels, the count of ports used for WaveMark channels and the count of event channels. You edit the channel settings by double-clicking on a channel, or by selecting a channel with the mouse or keyboard and clicking the **Edit** button. The **Reset** button deletes all editable channels and sets a standard sampling state.

Copy As Text, Log

These two buttons render the entire sampling configuration as text in a format suitable for inclusion in a laboratory notebook. **Copy As Text** places the text on the clipboard. **Log** copies the text to the Log window (so you can print it). The `SampleConfig$()` script command also give you access to this text with additional formatting opportunities.

Set channels

The **Set...** button opens a dialog in which you can set the maximum number of channels that can be stored in a data file created by the **File** menu **New** command. The standard maximum number of channels is 32, but you can create files with space for up to 400 channels. If you reduce the number of data channels and there were channels defined with numbers above the new limit, these higher numbered channels will be deleted.

Big file If you leave this box clear, data files are restricted to a size of 2 GB and are compatible with previous versions of Spike2. If you check this box, Spike2 will change the format of the data file so that up to 1 TB of data can be written to disk. If you create a *Big file*, it is incompatible with all versions of Spike2 before 6.11, and version 6.11 onwards will treat the file as read only. Spike2 version 7 can read and write such files.

There are two limits to the size of a Spike2 data file: the number of clock ticks in the file (limited to 2147483647) and the physical size of the file. The physical size before version 7 was limited to 2147483647 bytes, or 2 GB. These limits are related to the maximum size of a signed 32-bit number. For most users of Spike2, these limits are not a problem. However, it is possible to run out of disk space before you run out of clock ticks. For example, if you set the Waveform sample rate to be one clock tick and sample one channel, each sample point uses 2 bytes of data. If you set Burst mode sampling in the Resolution tab with n channels, each taking one sample per clock tick, then each clock tick is using $2n$ data bytes. This means that your maximum possible run time is limited to $2147483647/2n$ clock ticks.

Without using *Big file* mode, if you were sampling 32 channels of data at 31.25 kHz each (the maximum possible with a Power1401), in non-burst mode you would need to set the clock tick to 1 microsecond, which means that the clock tick limit would be some 35 minutes of sampling. However, the file size limit would be hit in around 17 minutes. If you changed to burst mode sampling, the clock could be run at 32 microseconds per tick, which extends the clock limit to 19 hours, but the file size limit is still 17 minutes.

If you check the *Big file* box, the maximum file size is increased by a factor of 512, and in theory, you can sample for the full time allowed by the clock tick period. However, if you sampled 32 channels at 31.25 kHz each in burst mode for the full time period, you would create a file of more than 128 GB. Such a file will be very difficult to backup and manipulate.

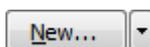
Unless you really cannot tolerate any gap in your data recording, you may find that it is better to use the Automation tab in the Sampling configuration and set the system to sample a sequence of smaller data files. These files will have a gap of a second or so between them, but will be much easier to backup. On the downside, this can make analysis more complicated as your data is spread over several files.

Once a file has been created as a Big file, it cannot be switched back. However, you can use the File menu Export command to export up to 2 GB of such a file to a new file without the big file format.

Previous versions Spike2 5.15 onwards can read files with up to 400 channels. Versions from 5.00 to 5.14 read files with up to 256 channels. Version 4.03 onwards reads files with up to 100 data channels. Versions before 4.03 can read files with up to 32 channels. You can use the File menu Export command to write channels from a data file to a new file with a suitable maximum channel limit so that it can be read by older versions of Spike2.

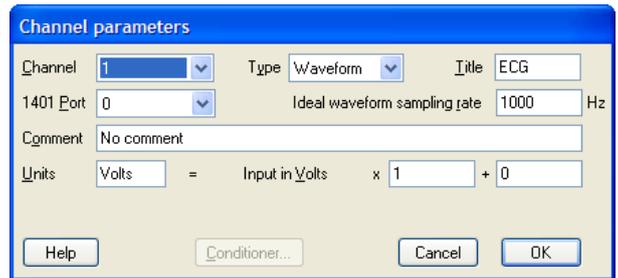
Create a new channel

The New... button creates a new channel at the next lowest free channel number. You can then edit the new channel to the desired state. To make several similar channels, set up the first channel, then click Duplicate (not for Talkers) to insert a new channel at the next free channel number (and physical 1401 port) with the same settings. You remove unwanted channels with the Delete button.



The small button next to the New... button is a short-cut to a list of channel types, and also gives you access to data from any installed Talker interfaces on your system. Talkers are discussed later in this chapter.

When you click the Edit... button or double-click on a channel, a new dialog opens where you set the channel characteristics. The first five fields apply to all data types.



Channel The channel number identifies this channel in the document. Channels 30 to 32 are reserved for Text markers, Digital markers and keyboard markers. The remaining channels are for waveform, WaveMark and event data. You can change a channel number, but not to a channel that is already in use.

Type This determines the type of the data to sample on the channel. You can set any of the following channel types:

- Off** The channel is unused, equivalent to deleting the channel
- Waveform** The channel holds waveform data
- Event-** Event channel timed on the falling edge of the data
- Event+** Event channel timed on the rising edge of the data
- Level** Event channel with the times of both data edges saved
- TextMark** The channel holds (short) text messages and their times.
- Marker** The data is an event with either a keyboard character or a digital value attached
- WaveMark** The data is a small waveform fragment, usually a spike shape

Title Up to 8 characters to identify the channel.

1401 port For a waveform or WaveMark channel, this is the 1401 ADC input number. For an event channel, this is the digital input port number (see page 4-5 for the connector pin numbers). Marker channels do not use this field. You will not be allowed to sample if you request a physical port number that does not exist in your 1401.

Comment Some text to give more information about this channel.

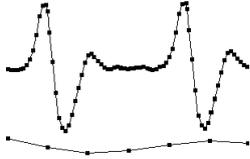
The **Conditioner...** button is present for Waveform and WaveMark data types and opens the signal conditioner dialog if a conditioner exists for the channel.

Maximum sampling rates

With a micro1401 or a 1401*plus*, the maximum continuous waveform and WaveMark throughput (sum of the rates of all channels) is 166 kHz. Power1401 and Micro1401 users can set any rate up to the maximum allowed by their 1401. If you use Event or marker channels or the Play waveform feature or the output sequencer at high data rates, you may have to reduce the overall sampling rate if you experience data overrun errors. You may find that display updates become slower as the overall sampling rate increases.

Of course, the rate you can achieve also depends on your computer and interface card. The (obsolete) ISA interface card is the slowest and will not achieve the maximum rates; upgrading to a PCI or USB interface will show significant throughput improvements. If you have a choice of USB or PCI interface on your 1401, USB 2 will provide by far the fastest throughput. USB 1 and PCI have similar data rates.

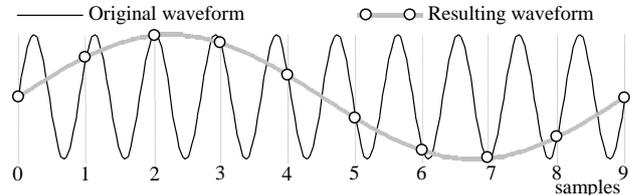
Waveform channels



The waveforms you record are continuously changing voltages. Spike2 stores waveforms as a list of numbers that represent the waveform amplitude at equally spaced time intervals. The process of converting a waveform into a number at a particular time is called *sampling*. The time between two samples is the *sample interval* and the reciprocal of this is the *sample rate*, which is the number of samples per second. The dots in the diagram represent samples; the lines show the original waveform.

Minimum sample rate

The waveform sample rate must be high enough to represent the data. The sample rate must be at least double the highest frequency contained in the data. If you do not sample fast enough, high frequency signals are aliased to lower frequencies. On the other hand, you want to sample at the lowest frequency possible, otherwise your disk will soon fill. Unlike many data capture systems, Spike2 lets you capture different waveform channels at different rates to minimise data file sizes.



Use of filters

Many users pass waveform data through amplifiers or signal conditioners with filter options such as the CED 1902 to limit the frequency range. Some transducers have a limited frequency response and require no filtering.

Input connections

Connect your waveform channels to the 1401 ADC inputs. Channels 0-7 (0-3 for a Micro1401) are the labelled BNC connectors.

Channel	8	9	10	11	12	13	14	15	Gnd
1401plus pin	1	2	3	4	5	6	7	8	9-15
Power (37-way)	28	29	30	31	32	33	34	35	1-19
Power (44-way)	33	34	35	36	37	38	39	40	3-10

Channels 8-15 are on the 1401plus 15-way front panel Cannon connector and on the Power1401 rear panel 37-way or 44-way connector. Pin numbers are given in the table. If you have an ADC expansion fitted you will have more channels; see the accompanying documentation for the connections. Power1401 top boxes can reassign the rear panel channels. The standard input Voltage range is ± 5 Volts. If you have a ± 10 Volt system check that the Voltage range is set correctly in the Edit menu Preferences.

Waveform dialog

The Waveform channel dialog has all the standard fields, plus:

Ideal rate Set the Ideal waveform sampling rate field to the desired sampling rate for this channel. The actual sample rate will be as close to this ideal rate as possible. You can see the actual rate in the Sampling Configuration dialog (in red if it is more than 20% different from the desired rate). If the rates differ too much, adjust the optimisation and clock settings in the Resolution tab.

Units This field holds the waveform units, up to 5 characters long. The following fields are the *scale* and *offset* to convert the input from Volts into user units.
 input in user units = input in Volts * *scale* + *offset*
 The *scale* is the number of units for every one Volt increase in input; *offset* is the value represented by 0 Volts at the 1401 input.

As an example, consider a situation where a waveform represents a position. 1 Volt is equivalent to 10 mm and 3 Volts is equivalent to 50 mm. In this case you would set:

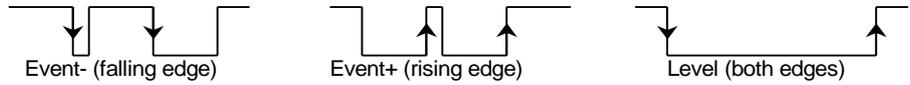
$$\begin{aligned} \text{scale} &= (50 - 10) / (3 - 1) = 20.0 \text{ mm V}^{-1} \\ \text{offset} &= 10 - (1 \text{ Volt}) * \text{scale} = -10.0 \text{ mm} \\ \text{Units} &= \text{mm} \end{aligned}$$

For the scaling to work as expected, the Edit menu Preferences option for Voltage range must be set correctly for your 1401. To display in metres in place of mm, set *scale* to 0.02, *offset* to -0.01 and *units* to m. The easiest way to set the calibration from known input data is to use the Analysis menu Calibrate command.

Event data



Spike2 stores time stamps very efficiently as integer multiples of the time resolution set in the **Resolution** tab of the sampling configuration. Simple time stamps with no other attached data are called *events*. Each event uses 4 bytes of storage. The 1401 recognises events as changes of state of TTL compatible signals connected to the 1401 digital input bits 15-8. We refer to these inputs as *event ports* 7-0. There are three types of event:



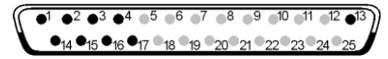
Event- Spike2 saves the time of the falling edge of the input signal. The minimum input pulse width is 1 μ s; wider is better.

Event+ The same as **Event-**, but Spike2 saves the time of the rising edge.

Level Spike2 saves the time of both edges. Pulses should be a minimum of 50 μ s wide or the time resolution set for sampling, whichever is the larger. Do not use this type unless you need the times of both edges.

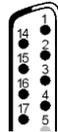
Micro1401 and Power1401 event port connections

Event ports 0 and 1 are BNC sockets on the front panel. If you have the optional Spike2 top box, event ports 7-2 are also on BNC sockets, otherwise you must use digital input bits 15-10. The digital input connector is on the rear panel. If you want to use the rear panel digital input connector for event ports 0 and 1 (to be pin compatible with the 1401*plus*) there is an option in the **Edit menu Preferences...** to use the rear panel connector for all events.



1401plus event port connections

In The event ports are bits 15-8 of the front panel digital input connector. Do not use the BNC sockets labelled Events 0 to 4; these have different functions: Event 1 is the digital marker input, 3 can be used to start sampling and also as a trigger for arbitrary waveforms. Events 0, 2 and 4 are used internally by Spike2 and have no external function.

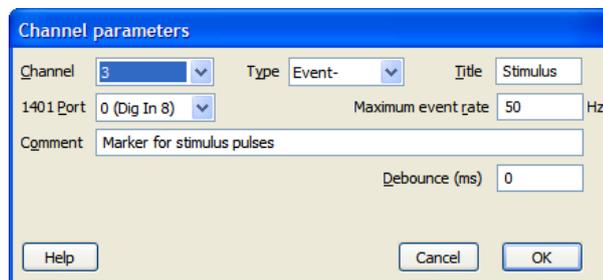


Input connections (digital input)

The pin connections for the digital input connector are the same for all 1401s:

Digital input bit	15	14	13	12	11	10	9	8	Gnd
Digital input pin	1	14	2	15	3	16	4	17	13
Event port	7	6	5	4	3	2	1	0	

Event dialog



The Maximum event rate is your estimate of the maximum mean event rate in Hz sustained over a few seconds. This is not the maximum instantaneous rate, which may be much higher. Spike2 uses this value to optimise buffer space allocation.

Debounce (ms)

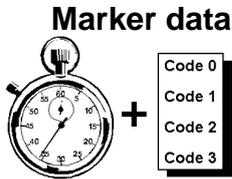
This sets the minimum period between events on this channel, useful if the signal comes from a noisy switch; events closer than this time are ignored. You can also set a negative value (-1) to keep all events, but adjust the times to be 1 clock tick apart. Set this field non-zero if it really is needed as it increases the processing required on every event.

TTL compatible signals

TTL stands for Transistor-Transistor Logic, a method for passing logical information between devices using voltage levels. Levels above 3.0 Volts are in the High state, levels below 0.8 Volts are in the Low state. Levels in between 0.8 and 3.0 Volts are undefined.

Do not subject 1401 TTL inputs to voltages above 5.0 Volts or less than 0.0 Volts. CED hardware has special circuits on TTL compatible inputs to provide some protection, however determined abuse will damage them. The 1401 TTL compatible inputs are

pulled up by a resistor to 5 Volts. They require a current of no more than 0.8 mA to pull them into the Low TTL state. Alternatively, you can connect them to ground to pull them low (this can be useful for the Event 3 input). See the *Owners handbook* of your interface for full details of all input ports.



Spike2 samples keyboard markers on channel 31 and digital markers on channel 32. A Marker is a 32 bit time plus 4 bytes of marker information. The first of these 4 bytes is the ASCII code of the keyboard character pressed by the user (channel 31) or an 8 bit digital code read by the 1401 (channel 32). The remaining three bytes are set to zero.

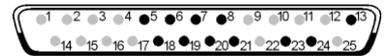
Spike2 treats all marker types identically once the data has been captured; they differ only in their source. You can also treat data types that are derived from markers (such as WaveMark, TextMark and RealMark) as if they were markers.

Keyboard markers

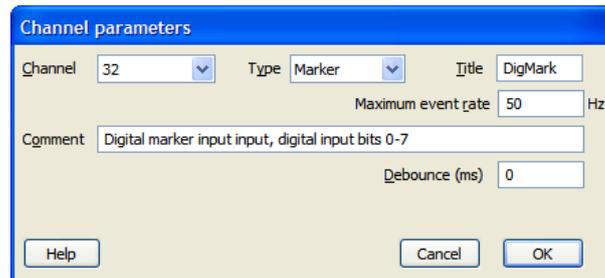
The timing of Keyboard markers is not precise; it depends on the load on the computer. Use the event inputs for exact timing. Any keyboard character that is not trapped for a special purpose (for example `Ctrl+L` opens the Evaluate windows) is recorded, but *only when the sampling document window is the current window*. In the special case where keyboard markers trigger data sampling (see page 4-18), the precise time of the trigger is stored. You can link keyboard markers to the output sequencer. The keyboard marker channel is always enabled.

Digital markers

The digital markers are timed as accurately as the event data. They record 8 separate channels of on/off information, or one channel of 8 bit numbers, or any combination in between. Digital marker data is sampled when a low going TTL compatible pulse is detected as described below. The data is read from bits 7-0 of the 1401 digital input.



Digital marker dialog



The Maximum event rate field determines the resources used to sample the channel. Set an estimate of the maximum sustained rate over several seconds. The Debounce (ms) field sets the minimum interval between events with the same code, useful if you have a noisy signal source.

Digital marker connections

The digital marker is read from data bits 7-0 of the digital input connector. The 1401plus needs a pulse on the E1 front panel input to flag a digital marker. The Micro1401 and Power1401 require a pulse on digital input pin 23. In addition to the data lines, there is an optional handshake (h/s) signal.

Marker data bit	7	6	5	4	3	2	1	0	Gnd
Digital input pin	5	18	6	19	7	20	8	21	13

To flag an event, apply a low going TTL pulse at least 1 μ s wide to the Event flag input. When the 1401 detects a falling edge at the Event flag input, it sets the h/s line active (within a few microseconds). The falling edge of the Event flag input latches the input data in the Micro1401 and the Power1401. The h/s returns to a non-active state after the 1401 reads the input. If you use a 1401plus you must keep the digital input data signals stable until the h/s line returns to the non-active state (this should never be more than 50 μ s after the Event flag input goes low).

Signal	plus	micro & Power
Event flag	E1	pin 23
h/s	pin 23	pin 24
h/s active	TTL high	TTL low

1401plus with the 1401-18 event discriminator

If your 1401plus has a 1401-18 event discriminator card (see the *1401-18 Programmable Discriminator* chapter), the 1401-18 uses all the pins on the digital input connector. The 1401plus duplicates the digital marker data bits on the same pin numbers on the digital output connector. However, pin 23 is not duplicated, so the h/s output for the digital marker is not available if you also have a discriminator card.

Output sequencer link

There are links between the digital marker channel and the output sequencer. If the REPORT instruction is used in an output sequence, this simulates a digital marker input pulse and causes the digital input to be read and the time to be recorded. As this is an internal activity, the handshaking described above is not available.

You can also use the MARK output sequence instruction to record a digital marker without reading the digital inputs (the instruction sets the 8-bit marker code). This instruction is often used to record output sequencer actions as part of the data file.

You can mix externally and internally generated digital markers, but this is not recommended unless care is taken to differentiate between the two sources of markers during analysis. This could be done by connecting the external marker handshake line to one of the digital marker data bits so that all external markers were flagged.

Warning: The DIBEQ, DIBNE, DIGIN and WAIT sequencer instructions use the same inputs as the digital marker and can cause digital marker events to be missed. Spike2 will warn you if this is an issue with your hardware.

Marker codes

When Spike2 displays markers, or data derived from markers, such as WaveMark or TextMark data, it shows the code of the first of the four markers. Marker codes occur in several other guises, for example to set trigger codes, as arbitrary waveform output codes and in the spike shape module.

Marker codes have values from 0 to 255. This is the same range of numbers that the 8-bit ASCII character set uses, and it is sometimes convenient to treat the codes as ASCII character codes (for instance when dealing with keyboard markers). At other times it is more convenient to deal with the codes as numbers.

Whenever Spike2 displays a marker code that is the same as the ASCII code of a printing character, it shows the printing character, otherwise it displays the character as a two digit hexadecimal code. Hexadecimal (base 16) numbers use the standard digits 0 to 9, but also use a to f (for decimal 10 to 15). Thus 00 to 09 hexadecimal is equivalent to 0 to 9 decimal. 0a to 0f is equivalent to 10 to 15 decimal. 10 to 1f hexadecimal is 16 to 31 decimal, 20 to 2f is 32 to 47 decimal and so on.

The printing characters are 20 to 7e hexadecimal, 32 to 126 decimal, as in the table. To find the hexadecimal code of a printing character, add the number above it to the number to the left of it. For example, the code for A is 41. To convert a code to a character, look up

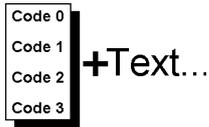
	+	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
20		!	"	#	\$	%	&	'	(*	+	,	-	.	/		
30		0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
40		@	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
50		P	Q	R	S	T	U	V	W	X	Y	Z	[\]	^	_
60		`	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o
70		p	q	r	s	t	u	v	w	x	y	z	{		}	~	

the first digit in the left column and the second in the top row. For example, 3f codes to ?, the intersection of the row for 30 and the column for f.

When typing marker codes (for example in the on-line Process dialog in Triggered mode, or when assigning codes in the spike shape module), type two hexadecimal digits for a code or type a single character to stand for itself.



TextMark data



This type is a combination of a marker and a text string. It is stored as a 32-bit time and 4 bytes of marker information followed by a text string that can be up to 80 characters long. This type allows you to insert timed comments into a data file. In the special case where the TextMark channel triggers data sampling, the precise trigger time is stored. From a script you can set longer or shorter strings with the `SampleTextMark()` command.

The channel number is forced to 30. If you select a channel number of 30, the data type is forced to TextMark. The Maximum sustained event rate is not used; please set it to a low number as it may be used in the future to optimise disk buffer sizes.

Serial line input

The Source field can be set to Manual or a choice of serial ports. The Terminator field sets the input character that marks the end of a text line. This field uses the same coding as for marker codes, so a single character stands for itself; two characters are interpreted as a hexadecimal code. Common codes are 0D for carriage return (the default) and 0A for line feed; 00 cannot be used. If no terminator is set, 0D is used.

With a serial port selected you can set the Baud rate, data bits, parity and handshaking. These must match the data source for reliable operation. See your computer hardware manual for pin connections and Baud rate limits.

When reading the serial line, characters codes below 32 are ignored unless they match the terminator character. The marker time is set when the first character arrives except for the special case where the TextMark channel triggers data sampling, when the precise trigger time is stored. The serial data can set the marker code by ending the text with a vertical bar followed by a decimal or hexadecimal number encoding the marker codes c_0-c_3 as $c_0+256*(c_1+256*(c_2+256*c_3))$. The vertical bar and number are not recorded. If `<term>` stands for the terminating character, the following are acceptable inputs:

```
Message from serial input, all marker codes will be 00<term>
Set marker code 0=0e (decimal 14), codes 1-3=00|14<term>
Set code 0=14, 1=02, 2=03, 3=40 as hexadecimal|0x4003020e<term>
```

Manual input

TextMark data is added to your file during sampling from a serial line and with the Sample menu Create a TextMark... command. You can also activate the dialog with the `Ctrl+T` key combination as long as the sampling data file is the active window.

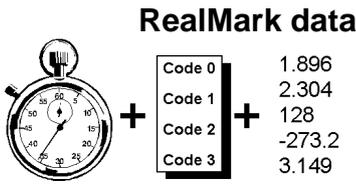


TextMark data is drawn as small rectangles. The rectangles are yellow unless the first marker code is non-zero, in which case the same colour coding as for WaveMark data is used.

Move the mouse pointer over a marker to see the attached text. Double click to view and edit the text and codes and display a list of the markers in the file.

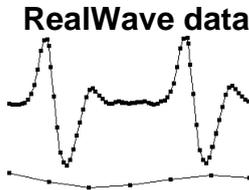
Time	C1	Text
2.55488	02	10 cmH2O
5.87008	03	Applied 10 ug/kg compound 1
8.14592	04	Stimulation amplitude 20
9.46432	04	Stimulation amplitude 40
13.11744	04	Stimulation amplitude 60
14.85568	02	12 cm H2O
20.20608	04	Stimulation amplitude 20
21.24544	04	Stimulation amplitude 40
27.03616	04	Stimulation amplitude 60
29.1328	02	14 cmH2O

If the TextMark channel is used, programmable signal conditioner changes are saved automatically as TextMark items.



RealMark data

This data type can be generated by external Talkers (not by a 1401). It stores a 32-bit time, 4 bytes of marker information and a user-defined number of single precision floating-point values. You can create and manipulate RealMark data from the script language or by processing measurements to a channel or through memory channels.

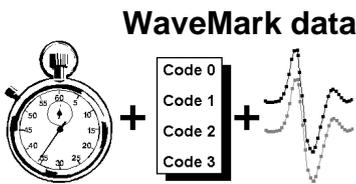


RealWave data

This data type is identical to waveform data except that the data is stored as 32-bit IEEE floating-point data, not as 16-bit integers. The channel has a scale and offset that is used to convert between RealWave and waveform data:

$$\text{RealWave data} = \text{integer data} * \text{scale} / 6553.6 + \text{offset}$$

RealWave can be sampled from external Talkers and created by scripts or as memory or virtual channels. Spike2 4.02 and older versions cannot open files holding this data type.



WaveMark data

This type combines waveform and marker data. It is stored as a 32-bit time and 4 marker bytes, followed by up to 126 waveform points on 1, 2 or 4 traces. The traces hold a spike shape. The first marker byte holds the spike classification code or 0 if it is unclassified. Script users can create WaveMark data for use off-line with up to 1000 data points.

Use WaveMark data where a high waveform sampling rate is needed to characterise very short events, for example nerve spikes. The 1401 searches the incoming waveform for sections that might include a spike. When the waveform crosses a trigger level, the signal is tracked to the next peak (or trough), and a data around the peak is saved.

WaveMark dialog

Channel parameters dialog box showing fields for Channel (5), Type (WaveMark), Title (Stereo), 1401 Port (2), Traces (1), Maximum event rate (100 Hz), Comment (No comment), Units (mV = Input in Volts x 1000 + 0), Points (32), Pre-trigger (10), WaveMark sample rate (20000 Hz), and buttons for Help, Conditioner..., Cancel, and OK.

The Maximum event rate is not the waveform sampling rate; it is the expected mean rate of events (these are usually spikes) per second.

The Units field is the same as for a waveform channel. You can use the Analysis menu Calibrate command to

calibrate known values. There are fields to set the number of data and pre-trigger points per event; these can be adjusted during template formation.

Points The waveform points per trace to store for each WaveMark on this channel. You should set this to the smallest value you can (the larger the value, the more space is used on disk, and the slower it is to process).

Traces A Micro1401 (but not the original micro1401) or Power1401 can sample multiple traces for stereotrode and tetrode data. Traces use consecutive ports; the 1401Port field sets the port for the first trace.

Pre-trigger The number of data points to keep before the first peak or trough to exceed the trigger level for this channel.

WaveMark sample rate This field sets the ideal sample rate for all WaveMark channels. Spike2 will adjust the sampling parameters to get as close to this rate as it can. See the Resolution Tab description for more information about rates.

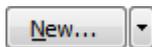
The total number of WaveMark traces you can sample depends on the 1401 type: 32 for Power1401, 16 for Micro1401, 8 for 1401plus and micro1401, 1 for 1401plus with the old non-silo ADC board. For example, a Power1401 could sample 4 WaveMark channels with 4 traces plus 4 with 2 traces plus 8 with one trace.

The Conditioner... button is enabled if a programmable signal conditioner is present.

Spike sorting Spike2 can match waveforms to a set of pre-determined shapes. This is normally used to extract single spike units from multi-unit recordings, but other uses are also possible (for example extracting R waves from ECG waveforms). If you record WaveMark data, Spike2 offers you a template setup window when you open a new file for sampling.

Talkers From Spike2 version 7.04, you can sample additional data channels from plug-in devices, called Talkers. A Talker provides one or more channels of data; the channels can be of any Spike2 data type. A talker can also provide additional keyboard marker channel input. You still need a 1401 interface to sample with Talkers.

Consider a blood pressure monitor that inflates a cuff periodically to measures systolic, diastolic and mean blood pressure. If this device has a computer interface, the output can be logged via a Talker. It could either provide three channels of RealMark data for the three outputs, or it a single RealMark channel with three attached data items.



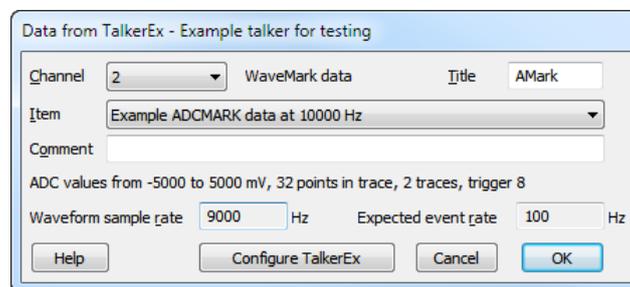
Channels from Talkers are added to Spike2 from the Sampling Configuration dialog by clicking on the small button to the right of the **New...** button. A drop-down list of possible channel types appears, with any Talkers known to the system at the bottom. If you select one of the normal channel types, this is equivalent to clicking the **New...** button, then choosing a channel type. However, if you select a Talker name, a new Talker dialog opens.



If any item in the drop-down list is disabled, this means that the available quota of channels for that channel type or Talker has been used. If a Talker is not currently connected to Spike2, the talker name is followed by an exclamation mark (!) to warn you that you cannot sample with it until it is connected.

For the purposes of this example we will assume that you have selected a Talker called **TalkerEx** that we provide with Spike2 as an example. You will find this Talker as **TalkerEx.exe** in the **Talkers** folder in the folder where **Spike2** is installed. Just double-click it to set it running, then select **TalkerEx** from the drop down list:

Typical Talker Dialog



When you select a Talker, a new dialog opens in which you can choose the data to write to a new channel. You should set the **Item** field first; this determines the type of channel to write to the data file and which additional fields to display in the dialog. If a Talker

supports multiple items, each item can be used to create a separate input channel. As a special case, a Talker can have 0 items, in which case the talker is assumed to provide additional input into the keyboard marker channel.

The **Channel**, **Title**, and **Comment** fields can be edited in exactly the same way as standard channel types sampled from a 1401 interface. The remaining fields are present to provide information and will change depending on the type of the channel.

The **Configure TalkerEx** button is enabled if the Talker is connected and if it supports configuration. It open a dialog that allows simple parameters of a Talker channel to be changed. Complex talkers will have specialised software drivers that provide their own, separate configuration.

Sample menu commands The Sample menu has a Talkers items that allows you to list all the Talkers that Spike2 knows about. From this menu you can get information about any connected talkers and also tell Spike2 to forget about talkers that you have previously used.

Talker documentation As Talkers are independent of Spike2, apart from the three example talkers described below, you will not find documentation for specific Talkers in the Spike2 documentation. A Talker that is written by CED to support specific hardware will have its own documentation, and we would hope that any Talkers provided by third parties to support their hardware will be suitably documented. If you have a problem with a non-CED originated Talker, it is unlikely that we will be able to provide you with support; you will need to take this up with the authors.

Spike2 keeps a list of all Talkers that it has been exposed to in a repository called `SP2TALKS.DAT` in the folder where Spike2 runs from. This allows Spike2 to configure sampling for talkers that are not currently connected.

Communications between the Talker and Spike2 are done using *named pipes*. This allows the Talker to run on the same computer as Spike2, or on a separate computer linked by a network. There is documentation and example code available to allow anyone who wishes to write their own Talker to connect to Spike2.

Example Talkers We include three example talkers with Spike2 in the `Talkers` folder in the same location where Spike is installed. You can run them on the same computer as Spike2 by double-clicking them. If you have a local area network connecting your computers, you can run them on a different computer by using a command line parameter:

```
C:\PathToTalker\TalkerName -sComputerName -n1
```

Where `C:\PathToTalker\` represents the location of the talker and `TalkerName` is the name of the talker. The `-s` option sets the name of the server computer that is running Spike2. If you omit the `-s` option, the server is assumed to be running on the same computer as the Talker. Talkers are identified by name and if you have multiple talkers, they must have different names. The `-n` option allows you to run multiple identical talkers; the number immediately after the `n` is appended to the Talker name, so `TalkerEx -n1` generates a talker called `TalkerEx1`. If the computer running Spike2 is called `Samuel` and your talker is in the folder `C:\Talkers`, you can connect with:

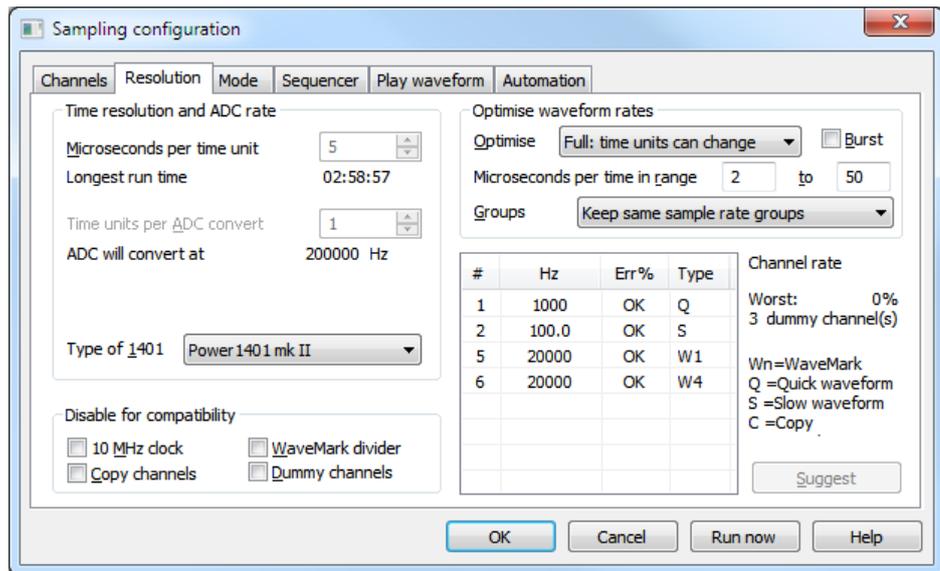
```
C:\Talkers\TalkerEx -sSamuel
```

The three example talkers are:

- `TalkerEx` Provides dummy channels of every type and generates dummy data. It also logs keypresses and sends them to the keyboard marker channel. This is purely an example; it has no practical use.
- `MKeys` Provides an additional marker channel that logs keypresses. This acts as a remote keyboard, so only makes sense if run on a different computer. You could use this to log responses from multiple subjects.
- `XKeys` Provides additional input to the keyboard marker channel. This acts like a remote keyboard, so only makes sense if run on a different computer. You could use this to log responses from a single subject.

Resolution The Resolution page of the sampling configuration dialog sets the time resolution of the document and the ADC sampling rate. To get the best possible results, you should read all this section. However, to get started quickly, follow these “cookbook” instructions:

1. Set the **Type of 1401** field as appropriate (set **General compatibility** if unsure).
2. Set **Optimise** to **Partial: fixed time units**.
3. Set **Groups** to **Keep same sample rate groups**.
4. Set **Microseconds per time unit** so that the **Longest run time** field is at least as long as the time you want to sample to the data file.
5. Clear all the enabled checkboxes in the **Disable for compatibility** section.
6. Uncheck the **Burst** checkbox (if visible) unless you want to run in burst mode.



This dialog controls how Spike2 optimises the sampling rates. Spike2 minimises the sum of the proportional errors between the desired and the actual rates for the waveform and WaveMark channels. By proportional we mean that an error of 200 Hz in a sampling rate of 10 kHz is the same as an error of 2 Hz in a rate of 100 Hz.

Microseconds per time unit This field sets the time units for a new data file in the range 1 to 10000 microseconds. All data of any type stored in the file occurs at a multiple of this time unit. The maximum time stored in a file is 2,147,483,647 units. The **Longest run time** is the maximum file length in days, hours, minutes and seconds for the current time units. To edit this field, set **Optimise** to **None** or **Partial**. If you select a **Power1401** or **Micro1401 mk II** or **-3** in the **Type of 1401** field, you can set this field to a resolution of 0.1 instead of 1 microsecond. 0.1 μ s resolution files are readable by Spike2 version 4.02 and later.

Time units per ADC convert This sets the ADC (Analogue to Digital Converter) clock interval in the units set by **Microseconds per time unit**. Each time the ADC is clocked either one sample is taken (non-burst mode) or a group of samples is taken (burst mode). To edit this field set **Optimise** to **None**. The **ADC will convert at** field is the equivalent rate in Hz for normal mode; it changes to **ADC burst rate** in burst mode.

Type of 1401 Members of the 1401 family have different capabilities. The choice you make here sets the absolute maximum settings for sampling. There is no guarantee that the maximum settings are achievable. The maximum rate depends on the entire sampling configuration, the input data load, the speed of the 1401, the speed of the interface connection (PCI, USB, ISA) and on the capabilities of the host computer. You can set:

<i>General compatibility</i>	Any 1401 except a 1401 <i>plus</i> with the old analogue card (an upgraded standard 1401). This is a <i>lowest common denominator</i> setting with a maximum waveform sampling rate of 166 kHz and does not assume that your monitor firmware is the most recent.
<i>1401plus, old ADC</i>	If your 1401 <i>plus</i> was upgraded from a standard 1401 without an upgrade of the analogue card you must select this option.
<i>micro1401/1401plus</i>	This setting is for a micro1401 or a 1401 <i>plus</i> with an up-to-date monitor. You can find if there is more recent firmware available in the Help menu About Spike2 command.
<i>Power1401</i> <i>Power1401 625 kHz</i>	These settings are for the Power1401. You can set <i>Microseconds per time unit</i> in units of 0.1, and <i>Time units per ADC convert</i> to 1 and sets the maximum ADC convert rate to 400 kHz (625 kHz with the later Power1401 625 kHz). Selecting any Power1401 or the Micro1401 mk II enables additional optimisations when calculating waveform rates (<i>10 MHz clock</i> , <i>WaveMark divider</i> , <i>Copy channels</i> and <i>Dummy channels</i>).
<i>Micro1401 mk II</i> <i>Micro1401-3</i>	Select one of these options to take advantage of all the capabilities of the Micro1401. This allows <i>Microseconds per time unit</i> to be set in units of 0.1, <i>Time units per ADC convert</i> to be set to 1 and sets the maximum ADC convert rate to 500 kHz.
<i>Power1401 mk II</i> <i>Power1401-3</i>	These settings allow <i>Microseconds per time unit</i> to be set in units of 0.1, <i>Time units per ADC convert</i> to be set to 1 and sets the maximum ADC convert rate to 1 MHz.
Disable for compatibility	If you replace a 1401 with a more recent model or upgrade Spike2 past version 6.05 you may get different waveform and WaveMark sampling rates due to new features that improve the match between requested and achieved sample rates. The new rates will be closer to the requested rates, but if you are art way through a study, compatibility with earlier data is more important than improved sampling efficiency. This part of the dialog disables new features to match the original rates. Fields that do not apply to the device selected in the <i>Type</i> of 1401 field are disabled. The settings of disabled fields are ignored. The fields are disabled for the 1401 <i>plus</i> and the original micro1401.
<i>10 MHz clock</i>	Check this box to force the <i>Microseconds per time unit</i> to be an integral number of microseconds. Before version 6.05 this control was available as part of the <i>Groups</i> field as <i>1 MHz, same sample rate groups</i> .
<i>WaveMark divider</i>	Previously, if you had WaveMark (spike shape) channels, their sampling rate set the maximum sample rate for all waveform channels. Now, we allow waveform channels to be sampled faster than WaveMark channels. Check the box to restore the old scheme.
<i>Copy channels</i>	If you sample the same channel as both a waveform and as a WaveMark, we only sample the data once for the WaveMark channel and take a copy for the waveform channel. Check the box to sample each channel separately.
<i>Dummy channels</i>	In old versions of Spike it was sometimes possible to get a better match to the desired sampling rates by adding more channels and ignoring their data. For example, if you requested 7 waveforms at 10 kHz each you got 9.99 kHz. Adding another waveform channel produced 10 kHz. Dummy channels do this automatically, and do not waste time moving data back to the host. Dummy channels are not used in burst mode.
Optimise	The <i>Optimise</i> field sets the parameters Spike2 changes to minimise sample-rate error:
<i>None: use manual settings</i>	You have control over the values in the <i>Microseconds per time unit</i> and <i>Time units per ADC convert</i> fields. In this mode click the <i>Suggest</i> button to change the fields to the values that minimise sample rate errors.

Partial: fixed time units You set the Microseconds per time unit field and Spike2 adjusts the Time units per ADC convert field to minimise the sampling rate errors. If there is more than one solution, Spike2 chooses the one with the slowest ADC convert rate.

Full: time units can change Spike2 sets the Microseconds per time unit and Time units per ADC convert fields. If all channels have a very slow rate, this can take a noticeable time. The Microseconds per time range fields set the acceptable range of time units. If there are multiple solutions, Spike2 chooses the one with the biggest Microseconds per time unit.

Groups This field places additional restrictions on how Spike2 maps the requested sample rates for all the waveform and WaveMark channels into achievable sampling patterns.

Version 3 compatible This gives the same waveform rates for a given Microseconds per time unit and Time units per ADC convert as version 3. Only use this if you upgrade from version 3 and it is vital that sampling rates match old data files.

Keep same sample rate groups If you select this option, Spike2 will make sure that all waveforms channels with the same ideal sampling rate have the same actual rate. **You will normally use this option.**

Ignore same sample rate groups This option may give smaller sampling rate errors. However, channels with the same ideal sampling rate may get different actual sampling rates. Some data analyses, such as waveform correlations or multiple channel averages and power spectra demand that channels have identical sampling rates.

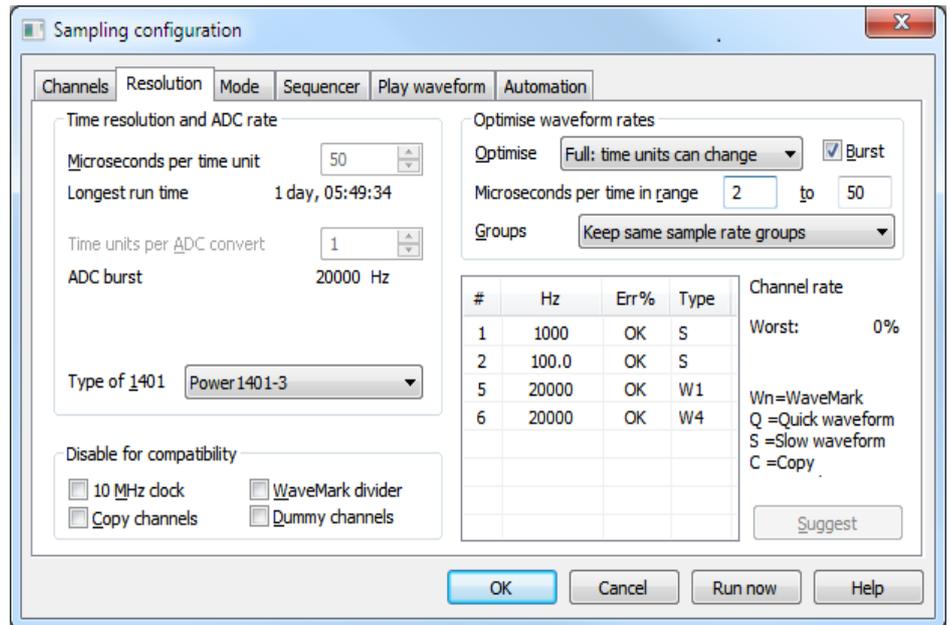
1 MHz, same sample rate groups This setting was present before version 6.05 but no longer exists. You can achieve this now by disabling the 10 MHz clock in the Compatibility section and selecting *Keep same sample rate groups*. If you read an old configuration that used this setting it will be translated to work in the same way.

Burst mode Burst mode sampling was added at version 6.06 and can be useful when you are sampling more than 1 waveform or WaveMark channel. If you have n waveform and WaveMark channels, the advantages of burst mode are:

- You may be able to run the Spike2 clock (set by the Microseconds per time unit field) n times slower. This can be useful if you need to sample for a very long time. However, you may run into problems with the size of the data file.
- You may be able to achieve sample rates per channel that are n times higher for the same Spike2 clock rate.

The disadvantage of burst mode is that the waveform channels will be sampled at times that are not exactly representable in terms of Spike2 clock ticks. Channels are sampled at precisely the correct interval, but each channel will be shifted sideways by up to half a Spike2 clock tick from when it was sampled. In many applications this will not matter, and the benefits of a longer run time or better sampling rate may outweigh this. Burst mode is most effective when you have a large number of waveform and WaveMark channels.

The **Burst mode** check box is visible if you have any type of Power1401 or a Micro1401 mk II or -3. If you check it, instead of sampling the ADC at equal intervals synchronised to the 1401 clock, it samples a burst of channels at a time, each burst being synchronised to the ADC clock. When you check the **Burst mode** box, the ADC will convert at field changes to ADC burst rate.



If you set the **Optimise** field to *None*, for full manual control, you will normally set the **Time units per ADC convert** field to 1 as higher values tend to reduce the benefits of burst mode.

Technical details

A master clock in the 1401 controls all sampling. The **Microseconds per time unit** field sets the tick period of this clock. The maximum sample time is 2,147,483,647 clock ticks. At 2 microseconds per tick this is 71½ minutes, at 10 this is almost 6 hours, and at 1000 this is nearly 25 days. All times in a Spike2 data file are multiples of this time unit.

The Analogue to Digital Converter (ADC) samples one input at a time and is shared between all the waveform and WaveMark (spike shape) input channels. The table to the right lists the maximum sampling rate in multi-channel mode for the ADCs in each 1401 when used in Spike2 and also if the

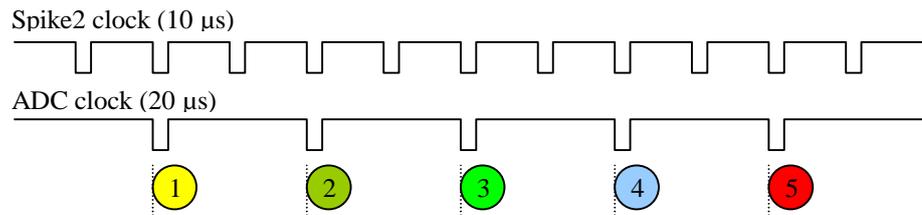
1401 type	Maximum rate	Burst
1401 <i>plus</i>	166 kHz	No
micro1401	166 kHz	No
Power1401	400 kHz	Yes
Micro1401 mk II/-3	500 kHz	Yes
Power1401 625	667 kHz	Yes
Power1401 mk II	1 MHz	Yes
Power1401-3	1 MHz	Yes

1401 supports **Burst** mode sampling. Using a slow interface (for example USB 1.0) will reduce the maximum rates.

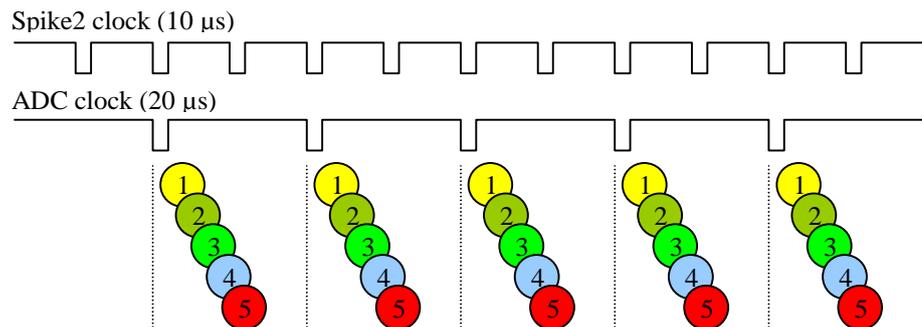
To illustrate the difference between non-burst mode and burst mode, we will consider what happens when we set **Microseconds per time unit** to 10 and **Time units per ADC convert** to 2 with five waveform channels in both non-burst mode and in burst mode. In the diagrams, the numbered circles represent the ADC sampling each channel and the horizontal position of the circle represents the time at which the sample occurs.

Non-burst mode

In non-burst mode, the **Time units per ADC convert** field sets how often the ADC samples in units of clock ticks. For example, with 10 microseconds per tick and the **Time units per ADC convert** field set to 2 (once every 20 microseconds), the ADC sample rate is 50 kHz. In a simple case with five waveform channels, the fastest each channel could be sampled is 10 kHz. Each sample is synchronised to the Spike2 clock and the pattern repeats every five samples.



Burst mode In burst mode, the Time units per ADC convert field sets how often a burst of ADC samples is taken. With the same settings of 10 microseconds per tick and the Time units per ADC convert field set to 2, the burst rate is 50 kHz. In a simple case with five waveform channels, five channels would be sampled in a burst, each channel sampled at 50 kHz. The interval between samples in a burst depends on the 1401 hardware and is typically the interval implied by the maximum sampling rate for the 1401 hardware.



As you can see, the ADC samples no longer fall exactly at Spike2 clock times. The times between samples on a channel are exact, but the entire channel may be displayed time shifted by up to half a Spike2 clock. In this case, channels 1 and 2 would probably be timed for the tick just before them, and channel 3, 4 and 5 for the next tick.

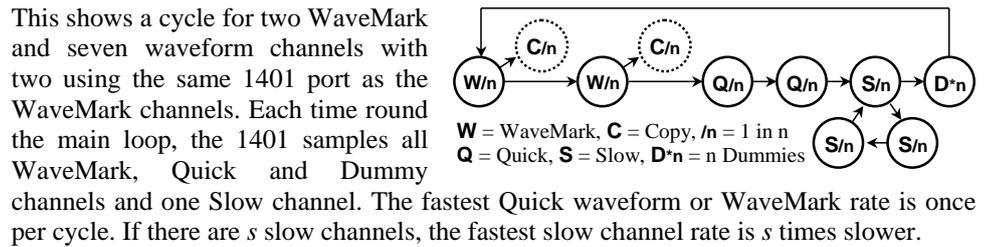
Cycle of channels To generate the sample rates, the ADC samples a cycle of channels. In burst mode all the channels in a cycle are sampled in a burst, in non-burst mode they are sampled one at a time. WaveMark channels are sampled once every time around the cycle. Some waveform channels are set as *Quick* and are also sampled every time round the cycle. The other waveform channels are set as *Slow*, and these share one position in the cycle. Quick and Slow channels save every n^{th} data point (n in the divisor in the range 1 to 2147483647). The Version 3 compatibility setting in Groups sets all waveforms as Slow channels and the maximum divisor to 65535.

With a Micro1401 mk II or -3 or any Power1401, there are further features:

WaveMark divider This enables down-sampling (taking 1 point in n) of WaveMark data. This allows you to sample Waveform channels faster than WaveMark channels.

Copy channel If you sample a waveform channel on the same 1401 port as a WaveMark channel, we can use the same data twice, once for the WaveMark channel and once for the waveform. These *Copy* channels behave exactly like a Quick channel, but are more efficient.

Dummy channel In non-burst mode, adding additional channels to the sampling loop sometimes gives a more accurate approximation to the desired sample rates. Adding a Quick channel would waste time transferring unwanted data to the host; a *Dummy* channel just throws the data away. Spike2 adds dummy channels automatically if they improve the channel sampling rates. Dummy channels are not added in burst mode.



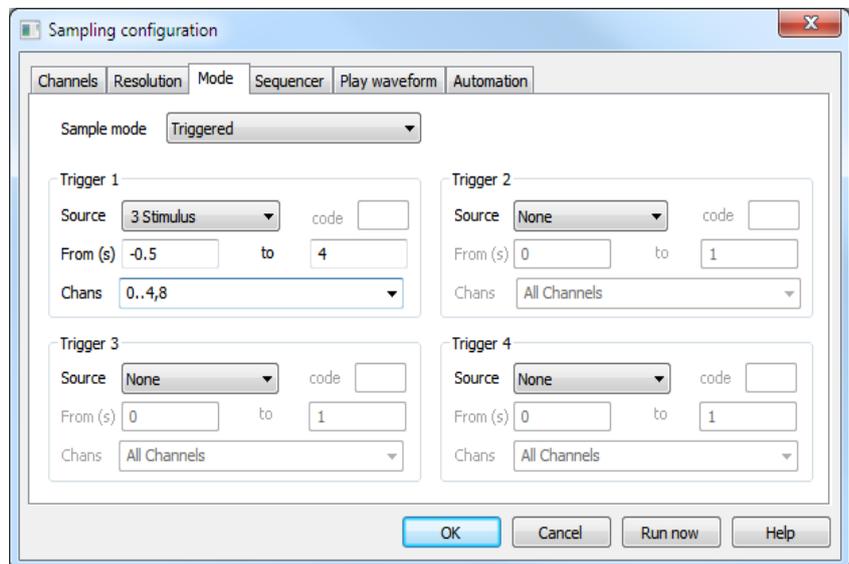
Spike2 searches all ADC rates allowed by the **Optimise** setting and the **Type** of 1401 field for the best combination of Quick, Slow, Copy and Dummy channels and the best value of n for each channel to get as close as possible to the ideal sample rates. With slow waveform rates there can be millions of combinations to search. If the **Channels** Tab feels very sluggish, set **Optimise** to **None**, set the channels, then restore the **Optimise** value. Impossible combinations display a warning in the lower left corner of the dialog.

If you check boxes in the *Disable for compatibility* section, this stops Spike2 taking advantage of optimisations that are not available for all 1401s. You might want to do this if you upgraded your 1401 and discovered that the sampling rates with your new 1401 were not the same as with the old one. Of course, the new rates would be closer to what you had asked for, but it might be important that they matched the old rates exactly.

The table to the left of the **Suggest** button lists the channels in order of descending sample rate error, showing the actual sample rate, the error as a percentage of the desired rate or OK if there is no error, and the channel type (WaveMark, Copy, Quick or Slow) and the number of Dummy channels.

Sampling mode

The **Mode** page of the sampling configuration dialog determines when data is captured by Spike2 and saved to the disk system. Whichever sampling mode you select, there is also a manual control during sampling that lets you disable data saving to the data document. Script users can enable and disable data saving channel by channel. In all modes, time passes at a constant rate, even when nothing is written to disk. When reviewed, areas of a file with no saved data are empty. There are three sampling modes: **Continuous**, **Timed** and **Triggered**.



Continuous mode

The simplest sampling mode is **Continuous** mode, which records data continuously.

Timed mode

In **Timed** data capture mode, data is saved to the document at intervals. You set the period for which data is saved and how often to save the data. Spike2 captures data on each channel in blocks. The start and end of a block does not, in general, fall at the same times as the start and end

of a sampling period. Spike2 saves *at least* the data you request, however you may get data before and data after the requested period. The times at which the blocks were requested are saved in the keyboard marker channel. Marker code 00 is placed at the start of each timed block and marker code 01 is placed at the end of each block.

Triggered mode

In **Triggered** capture there are four triggers. Each has an associated channel list. A trigger is an event or marker that causes data to be marked for writing to disk in a time range relative to the trigger. In triggered mode, any channel that is not associated with a trigger is recorded continuously. Each trigger has the following fields:

Source This should be set to **None** to disable the trigger or you can select a channel from the drop down list. You can trigger on any event, Marker or WaveMark channel in the sampling configuration.

Code If the trigger source is a Marker, WaveMark, RealMark or TextMark channel, you can choose to trigger only if the first marker code matches this field. You can set two hexadecimal digits to set the marker code or one printing character or leave the field blank for a trigger on all codes.

From This value is the offset from the trigger to the start of the area to record, in seconds. Negative values define pre-trigger start times, positive values start recording after the trigger. If you set too long a pre-trigger time, the data may have been discarded before the trigger is seen. You are warned when you start sampling if this is the case. Spike2 normally attempts to use an 8 MB data buffer, so this is not often a problem.

To This value is the offset from the trigger to the end of the area to record, in seconds. It must be greater than the time in the **From** field.

Channels Select **All Channels** from the drop down list, or type in a list of channels, for example 1..4,6,8 for channels 1, 2, 3, 4, 6 and 8. These are the channels that will be written to disk each time the trigger event is detected.

Data is always written to disk in complete buffers, so you will usually get more data written to disk than you request. Triggered sampling is usually used with fast waveform or WaveMark channels to save disk space in situations where only small sections of the data are interesting. If a new trigger occurs while data is being written, the trigger period is extended.

Special keyboard trigger features

If the user changes the state of the Write checkbox in the Sample control toolbar, a marker code is written to the Keyboard channel to indicate that writing to disk is enabled (code 00) or disabled (code 01). If the keyboard channel is used as a trigger, writing these codes does not trigger sampling.

Display during data capture

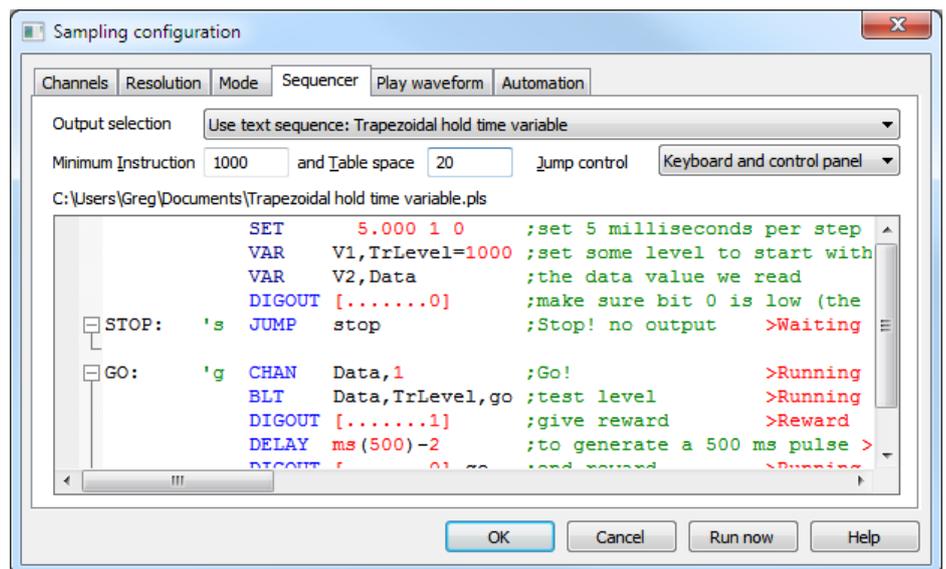
In all data capture modes, the on-line display shows newly sampled data, even when you are not saving to disk. Recent data is saved in a memory buffer in these modes so the current data is always available. However, if you scroll far enough back into an unsaved time region, the display may become blank.

Differences from previous versions of Spike2

Before version 5, there was only one trigger and WaveMark channels could not act as trigger sources. Triggered mode was handled by the 1401 and no pre-trigger data was available for display or writing to disk, nor could some channels be sampled continuously and other channels be triggered. If you use a triggered sampling configuration from a previous version, trigger 1 will be set to trigger all channels with **From** set to 0 and **To** set to the sweep length and the **Code** set to accept all marker codes. Because data is saved as complete buffers, you will probably save more data than in previous versions.

Sequencer

The Sequencer page of the Sampling configuration dialog sets the output sequence to use during sampling. The drop down list enables and disables sequence output and selects an output sequence file or sequences generated by the graphical sequence editor.



Output selection You can select None, Select a text sequence file, Use graphical editor, use the current file name, or you can select a recently used sequence file. The **Select a text sequence file** option opens a file dialog in which you choose the sequence file (*.pls) to attach to this sampling configuration. These files are created with the output sequence text editor or by exporting graphical editor sequences as text.

Minimum Instruction and Table space Unless you intend to change the sequence file during sampling you should set both these fields to 0. Spike2 reserves space in the 1401 for the list of sequencer instructions and for any table space that is used by the sequence set in the sampling configuration. However, if you change the sequence during sampling, the replacement sequences may need more space for instructions and a larger table space. You can use these two fields to reserve space; the actual space set is the larger of that required for the first sequence and the values you set in these fields.

Jump control The sequencer can be commanded to jump to specific places by keyboard commands, the sequencer control panel and always by the script language. You can limit the interactive options to prevent accidental jumps. Choose from: **Keyboard and sequencer control panel**, **Sequencer control panel**, **Script commands only**.

Text sequence open and edit If a file is selected it is displayed in the lower portion of the dialog. To modify the file, double-click in the displayed sequence to open it, then close the Sampling Configuration.

Graphical sequences If you select the graphical sequence editor, the lower portion of the dialog displays graphical sequencer control settings. See the next chapter for details.

Arbitrary waveform output

You can replay arbitrary waveforms during data capture. Up to ten different “Play wave” areas can be defined for output. Each area is identified by a key code, typically a printing character such as “A”. No two areas may have the same code. Script users can assign an additional code using `PlayWaveKey2$()`. You do not have to use a printing character, you can use a two digit hexadecimal code if you prefer, however codes 00, 01 and 02 are not allowed. The format of this code is the same as for marker codes (see page 4-7). Because you can trigger waveform output by recording keyboard markers with this code you should make sure that your code usage is compatible with key codes used in the output sequencer.

The waveforms are played from 1401 memory and are copied there just before sampling starts. This reduces the memory available for recording data. The maximum data you can store in the 1401 for replay is the free space in the 1401 less 256 kB which Spike2 reserves for recording. A micro1401 or a 1401*plus* or a Micro1401 mk II with 1 MB of memory can store around 750 kB of waveform. A Micro1401 mk II with a 2 MB memory can store around 1750 kB. A Micro1401-3 can store up to 3.7 MB. A 1401*plus* with expanded memory can store nearly 16 MB, a fully expanded Power1401 can store almost 256 MB or more, depending on the model. Script users can update the memory dynamically during replay, so huge memories are not necessarily required! The maximum size of a waveform area is 32 MB (assuming the 1401 has sufficient memory).

The final sample of the waveform area sets the output level after waveform output ends, so it is usually a good idea to make sure that the waveform output ends with a zero value.

The table shows the maximum rate we measured replaying two waveforms while sampling one. The input and output rates were the same. If you sample or replay more channels or use the output sequencer, the maximum rate will be lower. The rates were measured on a 450 MHz Pentium with a PCI interface card.

1401 kHz	<i>plus</i>	micro	Power	Micro
	62.5	100	250	167

The 1401 DAC output rate is derived by dividing down from a fixed frequency clock. This limits the sample rate resolution to 0.1 microseconds for Micro1401 and the Power1401 mk II and -3, 0.2 microseconds for the Power1401 and to 0.5 microseconds for the micro1401 and the 1401*plus*. These limitations can be significant, particularly if you replay imported data. For example, 44.1 and 48 kHz (often used in .WAV files for sound recording) cannot be exactly represented. You can use the `ChanSave()` script command to resample data to a different rate.

Play waveform toolbar

There is a dockable toolbar associated with the waveform output. This is enabled when you are sampling data with waveforms defined. The toolbar can be docked on any edge of the application and can be resized when it is floating. You can assign your own labels to the buttons, or you can let Spike2 generate labels itself of the form *Wave 0*, *Wave 1* and so on. The first button in the toolbar is used to stop a currently playing waveform.



Channels and DACs

Each play wave area contains from 1 to 4 data channels. You can select the 1401 DAC (Digital to Analogue Converter) each channel plays through. With multiple channels, all channels play at the same rate and all DACs update together. You can play data from waveform and WaveMark channels in a Spike2 data file, or data generated by a script. When data comes from a Spike2 file, the channels need not all have the same sample rate; Spike2 takes the rate of the first channel and interpolates data from the subsequent channels to make the rates the same. Script users can play arbitrarily long data by updating the waveforms in the 1401 online with the `PlayWaveCopy()` command. They can also modify the area sizes with `PlayWavePoints()`.

The rate at which a wave plays can be modified in the range 4 times slower to 4 times faster than normal (as long as your hardware can output fast enough). From a script, you can change the rate during sampling, even while the wave is playing.

Playing waves There are several ways to initiate output of a wave during data sampling:

- Click the associated button in the Play waveform toolbar
- Record the associated key code in the Keyboard marker channel
- Use the script language `SampleKey()` command to record the key code
- Use the output sequencer `WAVEGO` command

Unless you use the output sequencer to start the output, the associated key is recorded in the Keyboard channel and marks the time at which output was requested.

When a wave starts to play, there is a time delay of one waveform output between the moment that output is requested and the first data point appearing. This is not usually important, but at a slow replay rate, a delay of one sample could be significant. This delay can be useful if you are using the output sequencer as it gives the output sequencer the opportunity to know about a change in the DAC outputs *before* it happens.

Triggered output Each wave you play can be marked as “Triggered”. In triggered mode, when the waveform play is requested, output does not start immediately. Instead, the 1401 hardware waits for a trigger input (high to low edge) on the Trigger input for the Micro1401 and Power1401 (the 1401*plus* E3 input) unless this is routed to the rear panel Event connector pin 4 (Ground is pins 9-15) by the Edit menu Preferences. In triggered mode, the first data point is output at the time of the trigger.

If you use the output sequencer you choose between triggered and non-triggered in the `WAVEGO` instruction. You can also trigger the wave from the sequencer with the `WAVEST` instruction and detect if the wave has started to play with the `WAVEBR` instruction.

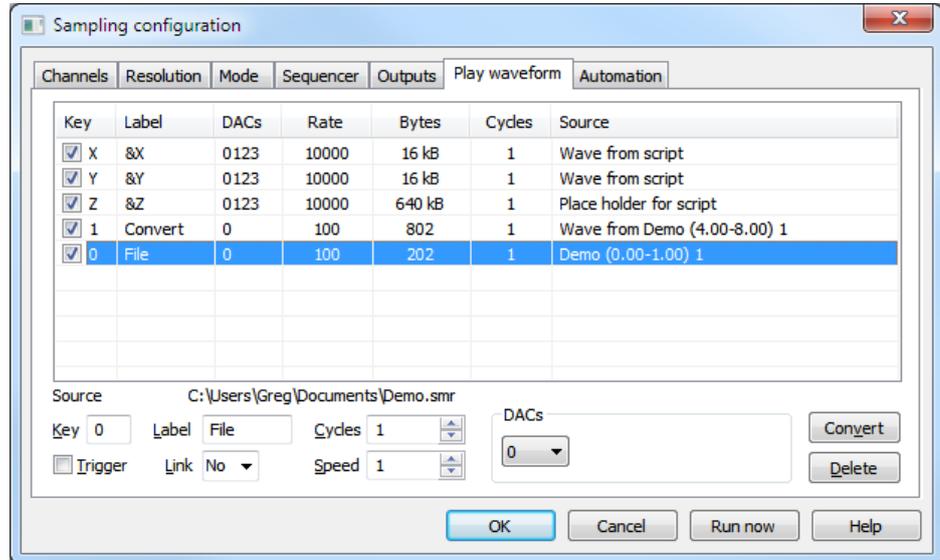
Repeated plays and links Each wave can be set to play cyclically for a set number of times. You can also link waves together if they have the same DAC output list. Linked waves play at the rate of the first wave, that is the sample rate of subsequent waves is ignored.

For example, you might have a sound output that needs to ramp up, stay at a constant level, then ramp down to zero. This could be done with three waves. The first holds the ramp up waveform, the second which repeats cyclically many times would hold the constant sound, and the final part would ramp down.

Scripts and the output sequencer can command a wave with many cycles to finish the current cycle then continue to the next linked area. You could use this to simulate a blood pressure signal with an occasional errant beat by having two waves, one with a normal beat set to repeat many times and one with the errant beat set to play once. By linking the two areas to each other, you get one errant beat after a fixed number of normal ones. Further, by using the randomisation functions in the output sequencer you could produce errant beats randomly and mark the times at which they occurred.

Play waveform

The Sampling configuration dialog Play waveform page holds the list of waves for on-line output. You add waves with the Sample menu Offline waveform output dialog and by scripts. The dialog settings are used the next time you sample data.



Key This is a checkbox plus a single character or two hexadecimal digit code in the same format as a marker code (see page 4-7) that identifies the wave; no other wave may have the same code nor may you use code 00. This code is added to the Keyboard marker channel when you click on a button in the Play waveform control bar.

If the checkbox is checked, space is reserved for the associated wave in the 1401 and if a wave is defined, it is copied to the 1401 before sampling starts. If you have defined a large number of data bytes for all the checked waves, it can take a significant time to copy them to the 1401.

Label A label of up to 7 characters that is used to label buttons in the Play waveform control bar. If you use &, the next letter is underlined on the button and you can use it as a short-cut to the key when the control bar has the input focus.

DACs This group of fields lists the Digital to Analogue Converters to play your waveform out of. You can change the list with the drop down lists inside the group. You cannot set two channels to use the same DAC output.

Rate The number of samples to output per second per channel. This is set when the wave is added to the list. You can vary the replay rate with the **Speed** control in the range 0.25 to 4.00. If the speed control is set to any value other than 1.00 the Rate field shows the multiplying factor as well as the rate.

Size This is the number of bytes of 1401 memory that are needed to hold the wave.

Cycles The times to play the wave. Use 0 for a very large number (about 4 billion).

Source This field describes where the data for the wave is stored. In the following descriptions, **Name** is a data file name, **sTime** and **eTime** are the start and end times of the data and **chans** is the list of channels to read the wave from.

Name (sTime-eTime) chans: the wave is in the data file. The full path to the file is displayed below the list of waves.

Wave from Name (sTime-eTime) chans: the wave is held in memory and saved in the sampling configuration and was read from the file. The full path to the file is displayed below the list of waves.

Wave from script: the wave is held in memory and saved in the sampling configuration and was generated by a script.

Place holder for script: a script reserved space, but did not generate a wave.

Trigger Check this box for waveform output that is enabled by play requests and that starts on a 1401*plus* E3 trigger or the Trigger input for other 1401s. If this is not checked, a play request starts the waveform playing immediately.

Link You can link waves with identical DAC channel lists together. Linked areas play in order with no time gap between them at the rate set by the first wave played.

Convert This button changes data held in a file into data held in memory and vice-versa.

Update Applies any changes you have made to the current wave.

Delete Delete the current wave.

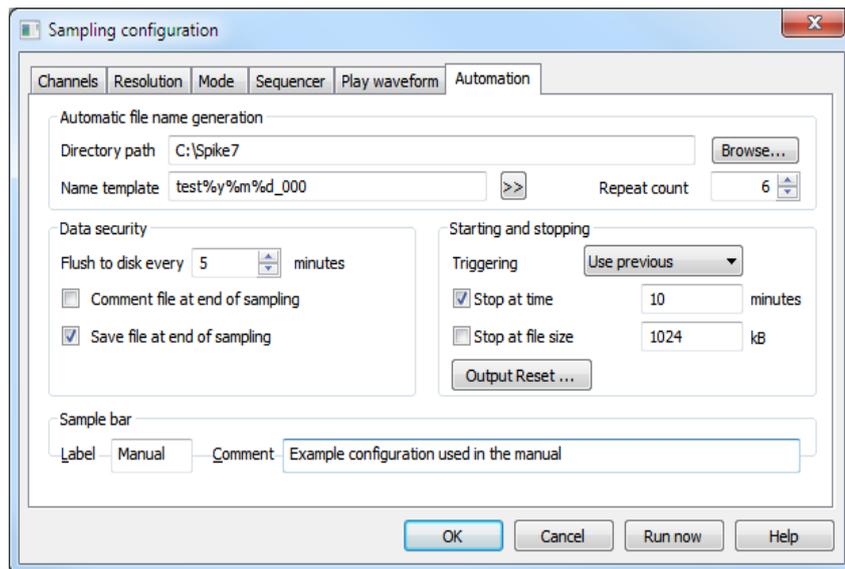
The line above the **Key** displays the path to any data file used as a waveform data source.

Automation

The Automation page of the sampling configuration dialog sets the file path and name for automatic data filing, sets the security parameters for new data files and restricts the total sampling time and the size of the data file. It also contains the label and comment used when you add the sampling configuration to the Sample Bar.

Automatic file name generation

Spike2 samples data to temporary files in the folder set by the Edit menu Preferences. These files have names like Data1, Data2 with no .smr extension. When sampling ends, you save the file to a name of your choice. You can automate file naming by setting a Name template (up to 23 characters), and Spike2 will generate a sequence of file names from it. If the Name template is blank, automatic name generation is disabled.



If the template does not end in a number, 000 is added before using it. To make a file name, Spike2 increments the number until a name is formed that is not in use in the Directory path folder (or the current folder if Directory path is blank). A template of test generates test000 to test999. A template of test10 generates test10 to test99. The Directory path must be less than 200 characters long. If the folders set by Directory path and in the Edit menu Preferences option are in the same disk volume, sampled files are renamed rather than copied to their destinations, which can save a lot of time, especially if you are using automatic sampling on a sequence of files.

With a file name template set, the generated name is used when the data file is saved. A different name can be specified using the File Save As... command. If there are no free names, or the path set for file saving does not exist, you are prompted for a file name.

Date and time specifiers

Within the file name, the character sequence % followed by one of y, m, d, H, M or S is converted into a two digit representation of the year, month, day, hours, minutes or seconds and %Y is replaced by the year as 4 digits. If the name ends with a date or time specifier, Spike2 adds _ to the name so that a date or time is not incremented. For example, test%y%m%d on December 19, 2007 becomes test071219_000. Click the >> button to insert a specifier at the caret or to replace the current selection.

Automatic sampling of a sequence of files

If you have set the Name template and Directory path fields, checked the Save file at end of sampling box and either or both of the Stop at time or Stop at file size boxes, you can enable automatic sampling of a sequence of files by setting the Repeat count field. Values of 0 or 1 sample a single file, larger numbers sample a sequence of files.

Sampling each file starts and stops based on the conditions in the **Starting and stopping** box. For each file except the last, when sampling stops the file is closed, and all result and XY views created by processing from it are closed. Sampling resumes with the next file in the sequence. The file sequence stops when:

- The number of files set in the **Repeat count** have been sampled
- There is no free disk space
- The name template cannot be incremented or there are no unused names
- There is an error during sampling
- The user clicks on **Stop** or **Abort** in the sampling control panel.

Data security For efficient sampling, Spike2 buffers several megabytes of the most recently sampled data in memory and writes data blocks to disk when the buffers are full. However if the power failed, this data would be lost. The **Flush to disk every n minutes** field sets how often the data buffered in memory is written to the physical disk to guarantee that your data is safe. There is a time penalty for doing this, so if you want the fastest possible sample rate you should turn this feature off (by setting a zero period). However, with it on, even if the computer power is lost or the computer crashes, your data should be safe up to at least the last flush time. You must run the `SONFix` program on such a data file to complete the data recovery and to tidy up data blocks written after the last flush.

If the **Comment file at end of sampling** box is checked, you will be prompted to provide a file comment when sampling finishes.

If the **Save file at end of sampling** box is checked, the new data file is saved to disk automatically when sampling finishes. If automatic filename generation is in use, the generated filename is used, otherwise the usual prompts for a file name are provided. This box must be checked if you want to sample a sequence of files automatically.

Sample Bar The Automation page also holds a label of up to 8 characters and a comment of up to 80 characters for the Sample bar. When configurations have been saved to a `.s2cx` or `.s2c` file, they can be added to the Sample bar by the **Sample** menu **Sample Bar List...** command and any label or comment is used to provide information about the configuration.

Once a configuration is in the Sample bar, you can open a new data file ready to sample with a click of the mouse.

Starting and Stopping The **Triggering** field presets the state of the sampling control panel **Trigger** checkbox:

- | | |
|---------------|--|
| Use previous | Uses the current state of the Trigger checkbox |
| Not Triggered | Forces sampling to start immediately; clears the Trigger checkbox |
| Triggered | The first sampled file is triggered, any repeats are untriggered |
| All triggered | The first file and all repeats are triggered |

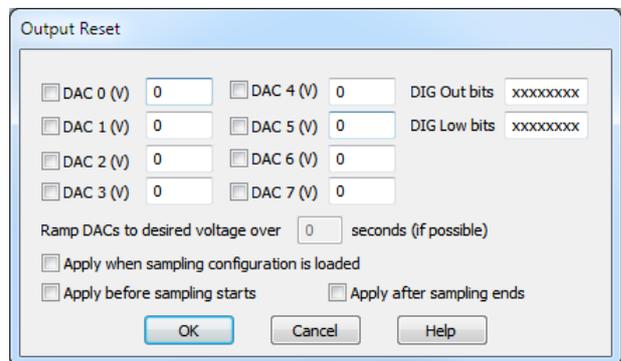
You can cause sampling to stop automatically at a set run time, or when the data file is a set size. If you do not check a box, the associated limit is not used. In general, you will get a little more data than implied by an active time or data limit as data buffered in the 1401 at the time the limit was reached will be added to the file. To sample a sequence of files automatically you must check at least one of **Stop at time** or **Stop at file size**.

Output Reset... The normal reset state of a 1401 is to have all DAC outputs at 0 Volts and all digital outputs set low. Unless you enable items in the Output Reset of Application Output Reset dialogs, Spike2 leaves all 1401 outputs alone except during sampling. Whatever state exists at the end of sampling will persist unless the 1401 is powered down. However, in some situations it can be important that the 1401 outputs are returned to a known state after sampling in case sampling was stopped early; this is where you would use the Output Reset dialog.

There are two sets of Output Reset information. One is associated with the Spike2 application, the other with the sampling configuration. The effect of the two sets of information is as if the applications settings were applied followed by the sampling configuration settings.

Feature	Application	Sampling Configuration
Opened from	Preferences, Sampling tab	Sampling configuration Automation tab
Dialog title	Application Output Reset	Output Reset
Data saved	In the registry	In the sampling configuration
Apply on load	When the program starts	When sampling configurations load
Priority	Values may be over-riden	Values are always used
Use when	Always	Set by the sampling configuration

The image to the right shows the dialog as opened from the Sampling Configuration Automation tab. When opened from the Edit menu Preferences, Sampling tab the dialog title changes to Application Output Reset and the Apply when sampling configuration is loaded check box becomes Apply when application starts.



DAC values Spike2 supports up to 8 DACs per 1401 interface. You can set the reset value (in Volts at the 1401 output) for each DAC. Only DACs that are checked will be reset.

Dig Out bits This field sets the reset values for bits 15-8 (from left to right) of the digital outputs as 8 characters. Each character can be 0, 1 or x standing for low, high and leave alone.

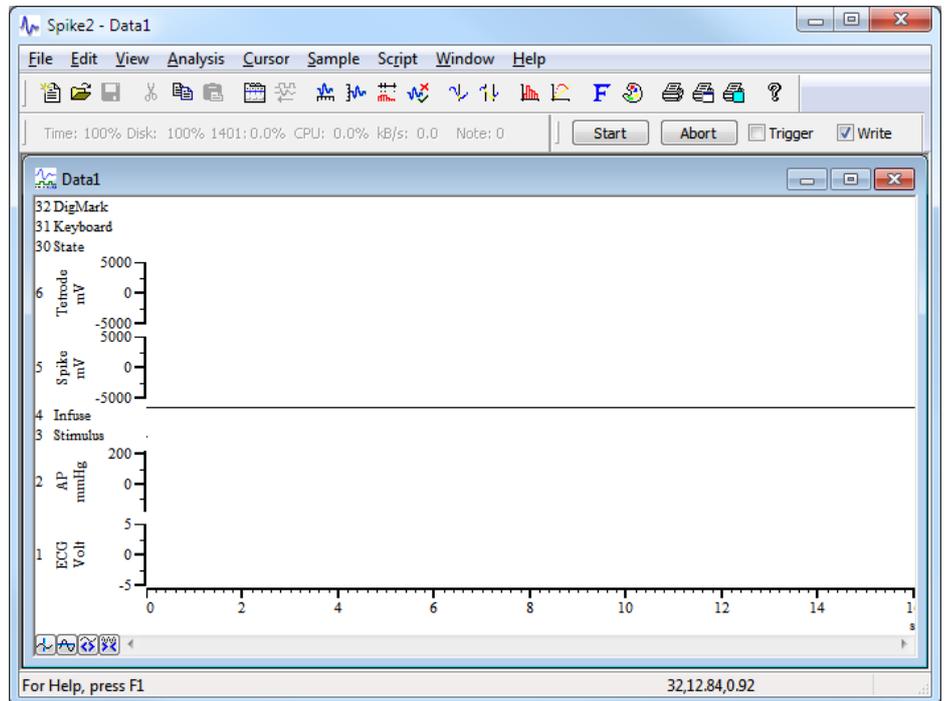
Dig Low bits This field sets the reset value for bits 7-0 (from left to right) of the digital outputs as 8 characters. Each character can be 0, 1 or x standing for low, high and leave alone.

Ramp DACs... Currently this field is inoperative and the DACs change to the reset value in one step. When implemented, the DACs will ramp from their current value to the reset value over the time period specified.

Apply... If you do not check any of the Apply... boxes, this dialog has no effect. You can choose when to apply the reset state: when the Output Reset information is loaded, before sampling starts, after sampling ends, or at any combination of these times.

Opening a new document

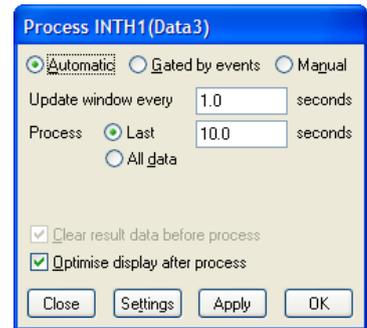
Once you have set the sampling configuration you can open a new data document. Select **New** from the **File** menu, then **Data Document** for the file type. Data documents differ from all other Spike2 documents as they are always stored on disk. Other document types are kept in memory until you save them. We keep data documents on disk because they can be very large. When you save a new data document after sampling, Spike2 moves it to the disk volume and directory you specify. When you use the **File** menu **New** command, Spike2 creates a temporary file in the directory specified in the **Edit** menu **Preferences**. If you do not specify a directory in the preferences, the location of the temporary file is system dependent.



The exact appearance varies, depending on the configuration. Sampling begins when you click **Start** in the Sample control toolbar (the **Sample** menu duplicates the controls in this window). If the **Trigger** box is checked, sampling waits for an external signal. You can set the display and analyses required before sampling. For example, to set an interval histogram you can select that analysis exactly as you did in the *Getting started* chapter. There is a difference, however. When you click on **New**, a new dialog appears.

Process dialog for a new file

You create result views with the Analysis menu New Result View command in the same way as when working off-line. However, the Process dialog, which controls when and how to update the result window, has additional options to work with a data file that grows in length. The radio buttons select Automatic, Gated by events or Manual updates (see the *Analysis menu* for a full description).



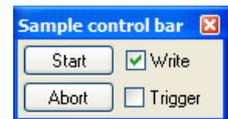
The standard on-line mode is Automatic. This mode adds new data to the result at a user-defined interval.

You can choose to accumulate a result for all the data, or produce a result for the last few seconds. The other two modes are for specialised uses and should not be used unless you are certain they are what you want. This dialog disappears once you select either OK or Cancel, however you can recall it with the Process command in the Analysis menu.

Sample control toolbar

The Sample control toolbar holds several buttons and a checkbox and controls the data sampling process. You can dock this bar to any edge of the Spike2 application window or leave it floating. The toolbar becomes visible (if it was invisible) whenever sampling starts unless the start command comes from the script language or the Edit Preferences Sampling tab option is used to prevent it. The position of the toolbar is saved in the sampling configuration. However, if the toolbar is visible and docked, we do not reposition it as we assume it is where the user wants it.

The Trigger field controls whether sampling starts immediately when you click Start or if it waits for a low-going TTL compatible trigger pulse. The Triggering field in the Automation tab of the sampling configuration dialog sets the initial Trigger state.



If you click Start with Trigger checked, Waiting flashes in place of the Start button until a suitable signal is applied to the Trigger input (Event 3 for 1401plus). Use this method to synchronise the start of sampling with an external event. Sampling starts within 1 or 2 microseconds of the external signal.



You can also decide which portions of your data are to be saved on disk, and which portions are of only transitory interest. The controls are duplicated in the Sample menu; however the Sample control toolbar needs fewer mouse clicks. The buttons are:



- Start** This is displayed before data capture starts. Click the button to start sampling. If Trigger is checked, Spike2 waits for the trigger before continuing.
- Stop** This is displayed while data is sampled. Click this button to stop sampling and keep the data. If no data was saved, the empty file is discarded.
- Abort** This button is used to abandon sampling and discard the new file. You can use this button before sampling starts, or while sampling is in progress.
- Reset** This button appears when you click Start. It stops sampling, discards any saved data, and waits for you to start sampling again with the same document.

The Write checkbox is normally checked, to save data to the data document. If you clear this box, no data is saved until it is checked again. Spike2 keeps a certain amount of data in buffers in memory, so it can display recent data on screen even if you have decided not to save it to disk. Whenever you change the state of the checkbox, a marker is added to the keyboard channel (code 00 when writing is enabled and code 01 when it is disabled). If you clear the box, the Write text moves from side to side to alert you and new data is drawn in the Not saving to disk colour.

High sampling rates

The code that transfers sampled data to disk runs in a separate thread with a high priority. This should ensure that data is saved unless another program creates a higher priority process that takes all the computer time. Buffer overflow can occur if the data rate is so high that the 1401 device driver cannot empty 1401 memory before it has become full. Spike2 detects buffer overflow and stops sampling if this happens.

If you suffer from buffer overflow problems with a Power1401 or a Micro1401 mk II or -3 interface, please check the following:

- Check that you have the latest 1401 firmware. The Help menu About Spike2 dialog box will tell you if more recent firmware is available.
- If you use USB to connect to your 1401 and your 1401 has a USB2 interface, make sure it is connected to a USB2 port on your computer.

If the 1401 detects that the host computer is slow writing to disk, it requests a “catch-up” mode where Spike2 abandons on-line display of new data (the display for new data is greyed out) to avoid competition between writes and reads in the disk system.

If the 1401 detects that the input event rate is too high for the 1401 to process, the special keyboard marker code `FF` is added to the keyboard channel. Sampling does not stop, as subsequent event times will be correct. If this happens, check that you have set realistic expected sample rates for the event channels.

The Sample Status bar

With a Power1401, or a Micro1401 mk II or -3 unit

Time: 100% Disk: 99% 1401: 4.3% CPU: 0.6% kB/s: 47 Note: 0

with up to date firmware, you can use the Sample Status bar to give you an indication of how hard the 1401 and the host PC are working to capture and transfer data to disk. If the bar is not visible during sampling you can show it from the Sample menu or by right-clicking in the toolbar area and selecting **Sample Status**. The bar shows five values:

- Time** Percentage of sampling time remaining before sampling stops. You can limit the sampling time in the Automation tab of the Sampling Configuration. If you do not limit the time, you can run for up to 2147483647 clock ticks (you can see how long this is in the Resolution tab of the Sampling configuration).
- Disk** Percentage of disk space remaining before sampling stops. You can limit the disk space in the Automation tab of the Sampling Configuration. If you do not limit the disk space, the limit is 1 TB in big file mode or 2 GB if not. The free space on disk is checked before sampling starts and the limits are reduced if there is less free space. Sampling stops if you get within a few megabytes of the free disk space when you started. If other activities are also filling the disk sampling may stop sooner, and possibly in a less controlled manner.
- 1401** The percentage of the time that the 1401 interface is using to transfer captured data back to the host computer. If this approaches 100% then it is likely that you are asking for more performance than your 1401 can achieve. This can also occur if you plug a 1401 with a USB 2 interface into a USB 1 port on your computer.
- CPU** The percentage of the available time that the data capture thread in the computer is using to process the incoming data from the 1401 and write it to the data file.
- kB/s** The number of kB (kilobytes) of sampled data transferred to the host per second.
- Note** The number of warning messages written to the Sampling Notes log. Double click this field to open the Notes logged window during sampling.

The 1401, CPU and kB/s fields display as *n/a* if you sample with a micro1401 or 1401*plus* or if the 1401 firmware is too old. You can download firmware updates from the CED web site. The CPU field will display as *n/a* if you have disabled the use of a separate thread for sampling (see the `Profile()` command in the script manual).

Saving configurations

It would be very tedious if you had to setup the exact screen configuration you wanted each time you sampled data. To avoid this, you can save and load sampling configurations from the **File** menu. The saved sampling configuration includes:

- The position and size of the application window
- The list of channels set for sampling and their sampling parameters
- The position of all windows associated with the new file
- The displayed channels and event display modes of the channels in time windows
- The name of any output sequence document to be used during sampling
- The list of waves and any associated waveform data for on-line waveform output
- The processing and update modes and positions of all result windows

The configuration does not include the contents of result windows. Whenever sampling finishes, the application saves the configuration as `last.s2c`. When you run Spike2, it searches for and loads the configuration file `default.s2c`. If this cannot be found, it uses `last.s2c`. These files are kept in the directory from which Spike2 was run. Remember that you can always recall the configuration that you used most recently, even if you forgot to save it.

If you have several configurations that you use very regularly, you can add them to the Sample bar with the **Sample** menu **Sample Bar List...** command. Once you have done this you will be able to start sampling with a saved configuration by clicking a button on the Sample bar (see the *Sample menu* chapter for a full description of the Sample bar). Alternatively, you could keep shortcuts to configuration files on your desktop and start Spike2 by double clicking them.

Warning We suggest that you do not rely on `last.s2c` for important sampling configurations as it is overwritten each time you sample. It is better to save your configuration to a named file and load it using the **File** menu or the **Sample Bar** or by double clicking the file.

Sequence of operations to set the configuration

This section describes a sequence of operations that you should follow to build a new sampling configuration from scratch. Once you have built a few configurations, it is often simpler to load an existing configuration and change the sections that do not fit your requirements, rather than re-build entirely. The steps are:

1. Open the **Sampling Configuration** dialog from the **Sample** menu.
2. In the **Channels** tab set the channels to be sampled and their sampling rates.
3. Adjust the **Resolution** values to give the best fit of sampling rates for any waveform channels.
4. Select the appropriate sampling **Mode** for your needs.
5. In the **Sequencer** tab select any required output sequence file or create a suitable graphical sequence or select **None**.
6. Set the **Automation** values as required by your application.
7. In the **Play waveform** tab select any waves needed for waveform output.
8. Click the **Run now** button.
9. Arrange the time view as you require and add any duplicate windows.
10. Add any result windows, set their update mode and position on screen.
11. Use the **File** menu **Save Configuration** command to save the configuration.

Once you have saved a configuration, you can re-use it by loading it before you use the **File** menu **New** command to start a new data file. You can combine the loading and opening a new file by adding the saved configuration to the Sample bar with the **Sample** menu **Sample Bar List...** command.

Data output during sampling

Overview

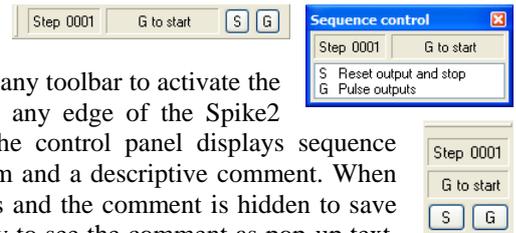
While sampling data you can generate precisely timed digital pulses and analogue voltages, monitor your experiment and respond to input data in real time with the Spike2 output sequencer. An output sequence is a list of up to 8191 (1023 for 1401*plus* and micro1401) instructions. The sequencer runs at a constant, user-defined rate of up to 100 instructions per millisecond with the Power1401 and the Micro1401 and up to 20 per ms with the micro1401 and 1401*plus*. The sequencer has the following features:

- It controls digital output bits 15-8 to produce precisely timed digital pulse sequences. In the Power1401 and Micro1401 it also controls digital output bits 7-0.
- It controls the 1401 DACs (Digital to Analogue Converters) to produce voltage pulses and ramps.
- It can play cosine waves at variable speed and amplitude through the DACs.
- It can test digital input bits 7-0 and branch on the result.
- It can record the digital input state or an 8-bit code to the digital marker channel.
- It supports loops and branches and can randomise delays and stimuli.
- It has 256 (64 in 1401*plus*) variables that can be read and set by on-line scripts.
- It supports a user-defined table of values for fast information transfer from a script.
- It can read the latest value from a waveform channel and the number of events from an event channel. Using this information, real-time (fractions of a millisecond) responses to input data changes are possible.
- It can control and monitor the arbitrary waveform output.

You write sequences with the text editor where each line of text generates one instruction or with the graphical editor where each graphical item generates one or more instructions.

Sequencer control panel

You can show and hide the control panel with the Sequencer Controls option in the Sample menu or by right clicking on any toolbar to activate the context menu. You can also dock it on any edge of the Spike2 application window. When undocked, the control panel displays sequence entry points as the key that activates them and a descriptive comment. When docked, the keys are displayed as buttons and the comment is hidden to save space. Move the mouse pointer over a key to see the comment as pop-up text. Click the mouse on a key and the sequencer will jump to the instruction associated with the key. The key is also stored as a keyboard marker. This is equivalent to pressing the same key in the time window or using the script language `SampleKey()` routine.



The control panel displays the next sequence step and any display string associated with it. You can use display strings to prompt the user for an action or to tell the user what the sequence is doing. Display strings are set with the text editor. The sequencer always starts at the first instruction. You cannot re-route the sequencer until sampling has started.

The control panel is always displayed if you start sampling from a Spike2 menu command or from a dialog. If the sample command comes from the script language or the Edit Preferences Sampling tab option is set to not display sampling toolbars, the control panel visible state does not change. The control panel position is saved in the sampling configuration. However, the position is restored only if the control is currently invisible or floating, and the saved position was floating; if the control panel is docked, we assume it is positioned where you want it.

Sequence jump disable

Sometimes you may want to stop users activating sequence sections with the keyboard or from the control panel. The Sequencer jumps controlled by field in the Sequencer tab of the Sampling configuration dialog lets you do this. The script language `SampleKey()` command can always activate sequencer sections.

Creating sequences

There are two ways to create output sequences: as an output sequence text file with the .PLS extension or as part of the sampling configuration using the graphical sequence editor. The table summarises the main differences between them.

	Text sequence	Graphical sequence
Edited with	Built-in text editor	Built-in graphical editor
Visualise output	No	Yes
Stored as	Output sequence .PLS files	Part of the sampling configuration
Implemented by	Machine code like language	Drag and drop editing
Ease of use	Takes time to learn	Very easy to learn and use
Flexibility	All features available	Uses pre-set building blocks
Timing	One instruction per text line	Several instructions per item

Graphical sequences are much easier to generate than text sequences. However, they have limitations and you can write more complex sequences using all sequencer features with the text editor. Sequences produced by the graphical editor are converted internally to the .PLS file format before they are used. You can save a graphical sequence as a .PLS file; you cannot convert a .PLS file into a graphical sequence.

The rest of this chapter describes the graphical output sequence editor, then the text editor, and finally the low level instructions used by both. You do not need to read about the low level instructions to use the graphical editor. However, knowledge of how the sequencer works will give you a better understanding of its capabilities and limitations.

Sequencer speed

The output sequencer runs at a rate set by a clock inside the 1401. You set the clock rate in milliseconds per tick in the Sequencer Tab of the Sampling configuration dialog when using the graphical editor or using the SET or SCLK directives in the text editor. The script command `SampleSeqClock()` can also change the rate. We allow intervals from 0.05 up to 3000 milliseconds on any 1401. The Power1401 and Micro1401 mk II or -3 allow intervals down to 0.01 milliseconds. If you set intervals less than a millisecond or use the sequencer text editor you should read the following information.

Sequencer technical information

The sequencer clock starts within a microsecond of recording time zero and is time locked to the 1401 event timing and waveform channel recording. Each clock tick books an interrupt to run the next sequencer instruction and updates digital output bits 15-8 if they were changed by the previous instruction.

An interrupt is a request to the 1401 processor to stop what it is doing at the earliest opportunity and do something else, then continue the original task. The time delay between the interrupt request and the instruction running depends on what the 1401 is doing when the clock ticks and the speed of the 1401. This delay is typically a few microseconds, so instructions do not occur precisely at the clock ticks but changes to digital output bits 15-8 do. Changes made by the sequencer to the 1401 DACs and digital output bits 7-0 occur a few microseconds after the clock tick.

The table shows the minimum clock interval, the timing resolution, the approximate time per step, the extra time used for cosine and ramp output and the time penalty for using the slow DIV and RECIPI instructions for each 1401. All values are in microseconds.

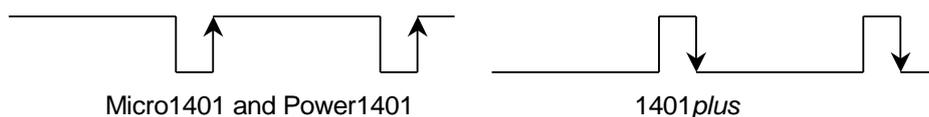
	Power	Power	Micro1401	micro1401	1401plus
Minimum tick	10	10	10	50	50
Resolution	1	1	1	4, 6, 10	4, 6, 10
Time used per tick	<1	<1	~1	<8	<10
Cosine penalty/tick	0.25	0.55	~1	~5	~10
Ramp penalty/tick	0.2	0.5	0.7	~3	No ramp
DIV, RECIPI penalty	<1	<1	<3	<10	<5

The **Minimum tick** is the shortest interval we allow you to set. The **Time used per tick** is how long it takes to process a typical instruction. The **Cosine penalty/tick** is the extra time taken per cosine output. The **Ramp penalty/tick** is the extra time taken per ramped DAC. Time used by the sequencer is time that is not available for sampling, spike sorting or arbitrary waveform output. To make best use of the capabilities of your 1401 you should set the slowest sequencer step rate that is fast enough for your purposes.

If you overload the system so much that the output sequencer cannot keep up, the result depends on the type of 1401. With the Micro1401 and Power1401, sampling stops with an explanatory message. The 1401*plus* does not detect this error and waits for the next clock tick to run the instruction.

The interval you set must be a multiple of the **Resolution** field for your 1401. This is not an issue for the Power1401 and Micro1401 mk II or -3. There is a choice of values for the micro1401 and 1401*plus* due to their hardware implementation. The table shows three possible values; in fact any value given by $n * m / 1000$ where n and m are integers greater than 2 is an acceptable resolution. Spike2 will not use a sequence if the interval is not an exact multiple of an achievable resolution for your 1401 as this would lead to inaccurate timing.

Sequencer clock output The output of the clock that controls the sequencer is available on the 1401*plus* OUT front panel connector and the Power1401 and Micro1401 Clock connector. This TTL output signal starts high with the Micro1401 and Power1401 and low with the 1401*plus*. The sequencer steps are synchronised with the rising edges of this output for the Micro1401 and Power1401 and the falling edges with the 1401*plus*. The minimum pulse width is 1 microsecond.



This clock runs whenever you sample with an output sequence. It can be used to synchronise external equipment. For example, if you connect this output to a counter, and place the counter in the field of view of a video camera that is used to record some other aspect of your experiment, the number on the counter links visual data with an exact time in the Spike2 file. If your sequencer ran at 1 millisecond per step, the counter would display time in milliseconds.

Marker channel and digital input conflict

If you record a digital marker channel, this uses the same hardware in your 1401 as the sequencer instructions `DIBEQ`, `DIBNE`, `WAIT` and `DIGIN` that read the digital input bits. Reading digital input bits 7-0 clears the hardware flag that indicates that a marker event has occurred. These instructions can cause events to be missed on the marker channel if they read the digital input at exactly the same time as the marker event arrives. The graphical editor commands that wait for a digital input condition also have this problem.

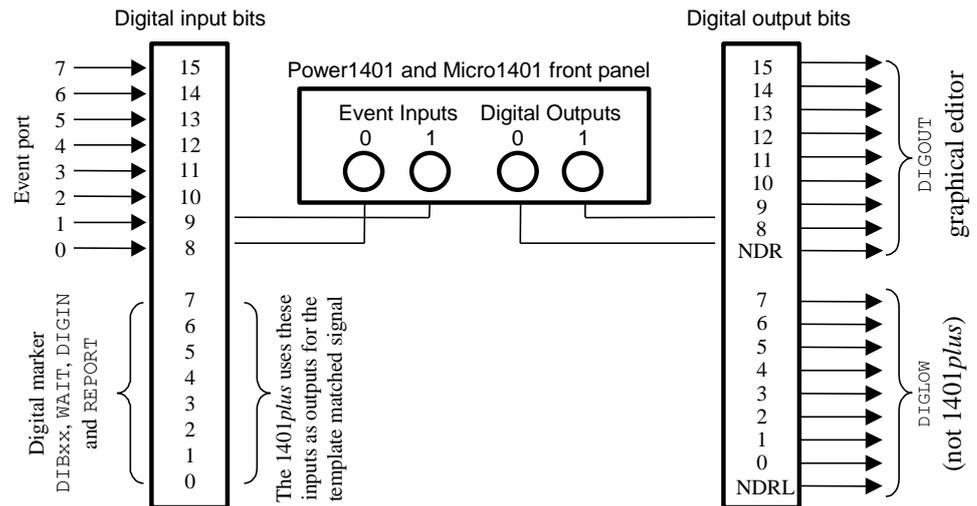
When you sample, Spike2 will warn you if there is a conflict in the use of these digital input bits. You can choose to hide the warning for the remainder of your Spike2 session. In October 2002 we released updates to the Micro1401 mk II and Power1401 firmware that allow reads of the digital input that do not clear the hardware flag. The warning is suppressed when the new firmware is installed in your 1401.

Digital input and output

The 1401 family has 16 digital inputs and 16 digital outputs. Digital input bits 15-8 are the event ports and are not used by the output sequencer. Digital input bits 7-0 are used for the Digital marker channel; you can also test these bits from the sequencer.

The sequencer controls digital output bits 15-8 individually to generate accurately timed pulses. Output bits 7-0 can be set with the DIGLOW instruction. The on-line template matching code can also use bits 7-0 of the output to signal spikes that match templates.

In a text-based sequence, the digital outputs are controlled by the DIGOUT and DIGLOW instructions and the digital inputs are tested with DIBxx, WAIT, DIGIN and REPORT. The graphical editor controls output bits 15-8 and tests input bits 7-0 with the delay and branching instructions. **Warning:** testing input bits 7-0 can conflict with recording events on the digital marker channel.

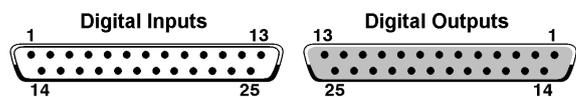


The digital input and output ports are 25-way connectors. The data and ground pins are the same on both. See your 1401 Owners Manual for full details of all pins.

Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	GND
Pin	1	14	2	15	3	16	4	17	5	18	6	19	7	20	8	21	13

Power1401 and Micro1401

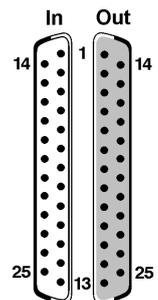
Digital output bits 8 and 9 are connected to the front panel as Digital Outputs 0 and 1 in addition to the rear panel. If you have the Spike2 top box, the remaining digital output bits 10-15 are also available on the front panel as Digital Outputs 2 to 7. The NDR output (New Data Ready, output connector pin 12) pulses low for 1 microsecond after a change to the digital output bits 8 to 15. NDRL (output connector pin 23) pulses low for 1 microsecond after a change to output bits 0 to 7.



1401plus

The 1401plus digital bits 7-0 are bi-directional (they can be set individually as inputs or outputs) and are connected to both the input and output connector. Spike2 normally uses them as inputs. The DIGLOW instruction has no effect on the 1401plus. There is no equivalent of the NDR signal available with the 1401plus.

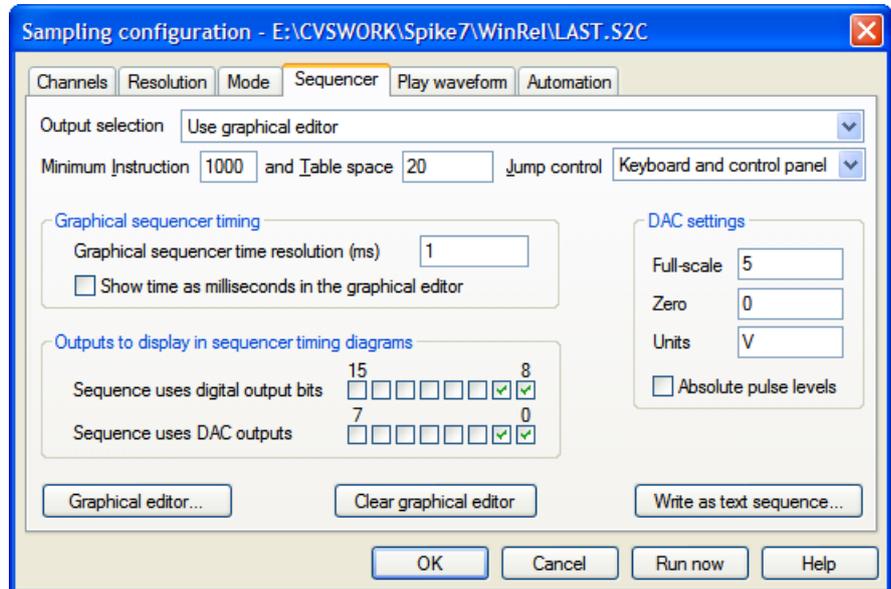
If you use the WaveMark *template matched* output signal, some or all of digital input bits 7-0 are set as outputs. The DIBxx, DIGIN and WAIT instructions read back the output state of the bits used for template matching. Do not drive these pins as inputs when they are used as outputs!



If the 1401plus has a 1401-18 event discriminator card, digital input bits 7-0 are the *event detected* TTL output connections; use the output connector bits 7-0 for the input signals.

The graphical editor

To view the graphical editor, open the Sampling configuration dialog and select the Sequencer tab. Then select **Use graphical editor** from the Output selection drop down list. The editable fields in this dialog set values that apply to the entire sequence:



Minimum Instruction and Table space

Unless you intend to replace the graphical sequence with a text sequence during sampling you should set both these fields to 0. Spike2 reserves space in the 1401 for the list of sequencer instructions and for any table space that is used by the sequence set in the sampling configuration. However, if you change the sequence during sampling, the replacement sequences may need more space for instructions and a larger table space. You can use these two fields to reserve space; the actual space set is the larger of that required for the initial sequence and the values you set in these fields.

Jump control

This field applies to both graphical and text sequences. It allows you to stop interactive control of sequence jumps to prevent accidental changes caused by a user keypress or mouse click in the sequence control panel. The `SampleKey()` script command can always cause the sequence to jump. You can choose from: **Keyboard and sequencer control panel**, **Sequencer control panel** and **Script commands only**. The script language equivalent is `SampleSeqCtrl()`.

Sequencer time resolution

This sets the time resolution of your sequence and the clock interval of the sequencer clock. This is also the minimum duration of any pulse. All actions in the sequence occur at integer multiples of the time you set here. You can set values in the range 0.01 to 3000 milliseconds. You need a Power1401 or Micro1401 mk II or -3 to set values less than 0.05 milliseconds (see page 5-2). Some actions take more than one clock interval.

Show time as milliseconds

Check this box to display and edit time in the graphical editor as milliseconds and not seconds. This is purely for your convenience; if your sequence sections are all less than a second you will probably find it more convenient to use milliseconds.

Sequence uses digital output bits

Check the boxes for the dedicated digital outputs that you will use. Only these outputs will appear in the editor, reducing visual clutter. If you do not require any digital outputs, clearing all the check boxes will save an instruction at the start of each sequencer section.

Sequence uses DAC outputs

Check the boxes for the Digital to Analogue Converters (voltage output devices) that you will use in your sequence. Unused DACs are not included in the graphical editor (to reduce visual clutter) and no sequence code is generated for them, which saves instructions at the start of each sequence section.

If you use one or more DACs for arbitrary waveform output only do not include them here unless you want to be certain that they have a defined value when you enter a sequence section. All 1401s have at least DACs 0 and 1; the 1401*plus* and the Power1401 have DACs 0 to 3. Top boxes allowing up to 8 DACs can be added to recent 1401s.

DAC full-scale, zero and units You can define the DAC outputs in units of your choice. 1401 DACs normally have a range of ± 5 Volts, but ± 10 Volts systems exist. Set **Units** to the units you want to use. Set **Full-scale** to the value in these units that corresponds to the maximum DAC output. Set **Zero** to the value in your units that corresponds to a DAC output of 0 Volts.

For a ± 5 Volt system calibrated in Volts, set **Full-scale** to 5, **Zero** to 0 and **Units** to v. If you want the output in millivolts, set **Full-scale** to 5000, **Zero** to 0 and **Units** to mV.

Absolute pulse levels The DAC pulses take their starting level as the current DAC value at the pulse start time. The DAC then changes to another value, then back to the original level. Normally you define pulses in terms of the pulse amplitude relative to the starting level and all pulses add. If you check this box, then you set the absolute level for the DAC to change to.

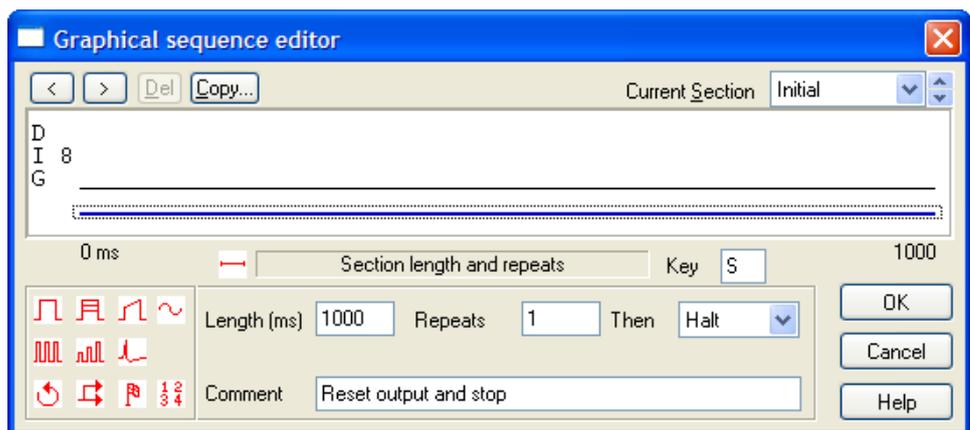
Write as text sequence... You can use this button to output the graphical sequence as a text sequence. You might do this if the graphical interface almost does what you need and you need to hand-edit a few extra instructions, or if you need to save the sequence for documentation purposes, or if you want to generate several sequences to use in a sampling session so that you can change them while running. A file selector dialog opens for you to choose a name for the .PLS file.

Clear graphical editor... This wipes all graphical sequences. You must confirm this action as you cannot undo it.

Graphical editor... This opens the editor so you can make changes to the sequence.

Getting started To get accustomed to the graphical editor, we will produce 10 millisecond wide TTL pulses at 1-second intervals from digital output bit 8. We will associate the key G for Go with the pulses, and the key S for Stop will stop them.

Open the Sampling configuration dialog **Sequencer** tab and select **Use graphical editor**. Click the **Clear graphical editor** button to remove any previous sequence. Set the **Sequencer** time resolution to 10 milliseconds and check the **Show time as milliseconds** box. We do not need any DAC outputs, so clear all the DAC outputs check boxes. We only need one digital output, so check the digital output 8 box and clear all the others. Now click the **Graphical editor** button.



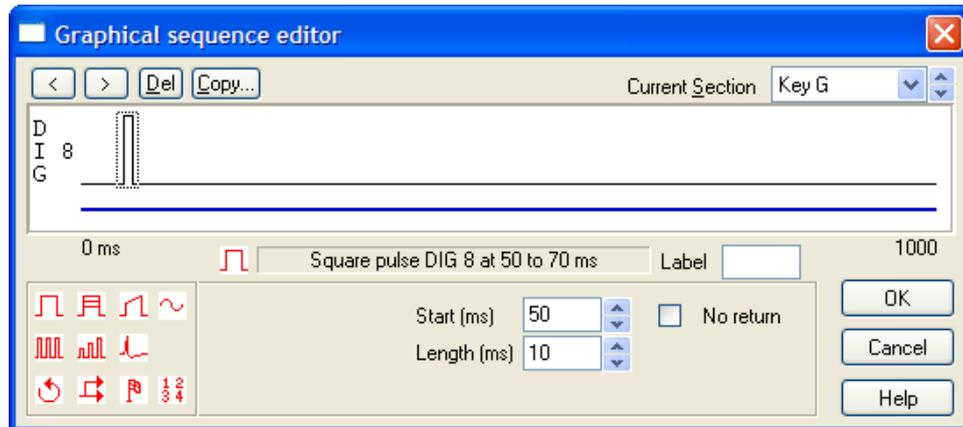
At the top right of the new dialog there is a drop down list to select the current sequence section. A section has a length and repeat count. When the repeats are done, the section either stops or jumps to another section. To start with, make sure **Initial** is selected as the **Current** section. The **Initial** section always runs first and it sets the initial conditions.

The graphical representation of a section always contains a *control track* drawn as a thick line at the bottom. We choose output on digital bit 8 only, so there is a single digital output in the remainder of the space. There is always one item selected in this area; the selected item has a grey rectangle around it. Click on the control track now.

In addition to setting the initial state, we will use this section to stop our pulse output, so set the **Key** field to **S**, and set the **Comment** field to **Reset output and stop**. This comment will appear in the sequencer control panel for the **S** key. You can leave the **Length**, **Repeats** and **Then** fields at their default values (1000, 1 and **Halt**).

The next task is to create the pulse train and associate it with the **G** key. Select **Key G** in the **Current section** field. The control track will be selected; click on it if it is not selected. Set the **Key** field to **G**. Set **Repeats** to 0 to mean repeat forever and the section length to 1000 milliseconds. The **Then** field will grey out as the repeats never end. Set the comment to **Pulse outputs**.

There is a palette at the bottom left of the dialog; you can drag icons from the palette and drop them on the DAC, digital output and control tracks. Click on the top left item in the palette and drag it to the digital output track and release it to create a pulse. Edit the start time of the pulse to 50 and the length to 10 milliseconds.



This completes the pulse setup. Click the **OK** button to return to the Sampling configuration dialog. Click the **Run now** button to start sampling with the output sequence you have just created. The sequencer control panel will display two items that you can select: **G Pulse outputs** and **S Reset output and stop**.

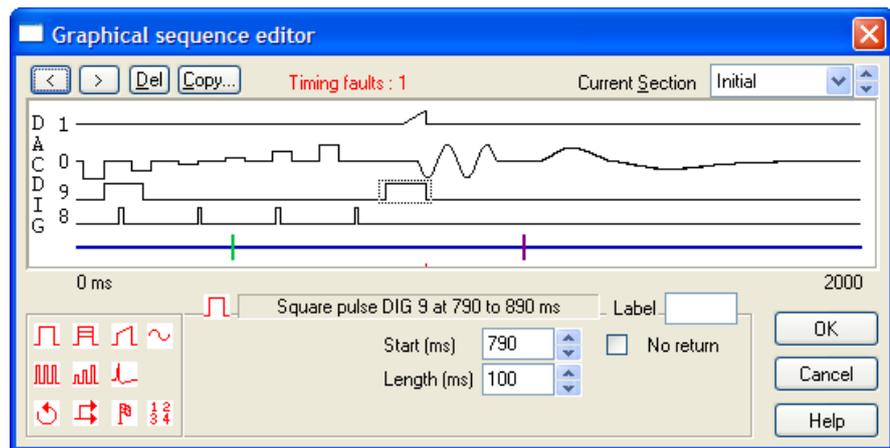
Start sampling – the sequencer will run the **Initial** section, which sets the starting values for all the digital and DAC outputs we have chosen to control. We did not request any other action in this section so nothing will happen until you use the **G** button in the sequencer control panel or select the sampling window and press the **G** key to start the output pulses from digital output 8. You can stop the pulses with the **S** button in the control panel or the **S** key on the keyboard.

Graphical editing

To open the graphical editor, select **Use graphical editor** and click the **Graphical editor...** button in the **Sequencer** tab of the **Sampling configuration** dialog. You can resize the dialog by clicking and dragging an edge. Double-click the title bar to maximise the dialog, double-click again to minimise it. The editor window has five areas:

1. a window with a graphical representation of the output sequence
2. above the graphical window are controls to iterate through and delete selected items, a message area and the **Current section** selector
3. the lower left-hand corner holds a palette of items to drag into the graphical window
4. the lower right hand corner has control buttons
5. the settings for a selected graphical item lie between the palette and the buttons

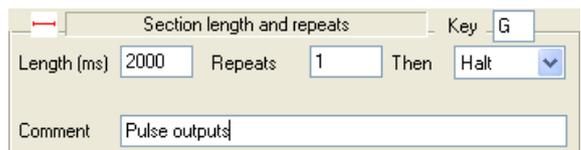
The **OK** and **Cancel** buttons both close the dialog. **OK** accepts all changes, **Cancel** rejects all changes. The **Help** button displays the information you are reading!



Sections There are 27 sections: **Initial** and **Key A** through **Key Z**. The **Current section** field sets the section to display and edit. The **Initial** section runs when the sequence starts; in some cases this may be the only section you need. The remaining sections are optional. A section displays a representation of the output for each DAC and digital output that is set for use in the **Sequencer** tab of the sampling configuration. There is a thicker line at the bottom for the control track, which holds all other sequencing actions.

Selecting items Click an item in the graphical view to select it. A grey rectangle marks the selected item and the item settings appear below the display. The **<** and **>** buttons at the top left select the previous and next item on the current track; they are very useful when items overlap.

Section settings Click on the control track clear of any dropped items. You can now set the associated key, the section length, number of repeats, action when the repeats are done, and a comment that is displayed in the sequencer control panel to identify sections with keys. You can set a section length up to 10000 seconds and up to 1000000 repeats! If you want a section to repeat forever, set the repeat count to 0. The **Then** field determines what happens when the section ends. You can stop the sequencer with **Halt**, or select another section to run.

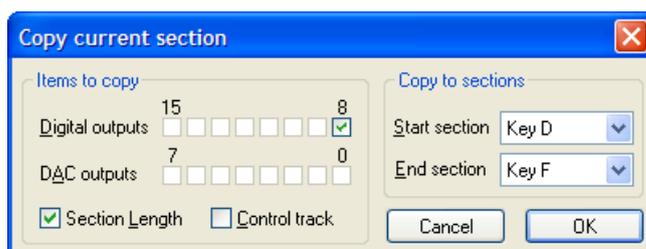


Key field Normally, the **Key A** to **Key Z** section are assigned A to Z and the **Initial** key is blank. However, you can set the key to any of A-Z, a-z or 0-9 (upper and lower case are different). If you delete the key, you cannot run the section from the sequencer control panel. Two sections that contain outputs cannot share the same key.

Adding and deleting items The graphical palette at the bottom left-hand corner of the dialog contains all the items that you can add to the display. Move the mouse over an item so see a short description. Click an item and drag it to a suitable track, then release to add it to the section. The Del button on the top line of the dialog removes the selected item. You cannot remove the control track or the lines that represent the initial state of the DACs and digital outputs.

Timing faults The sequencer attempts to match the timing you request. If this cannot be done, timing conflicts are marked in red below the control track and the number of conflicts is given in the message area. You will usually get a conflict if you try to position any action at the start of a section. This is because the first instructions in a section set the initial digital output state and the state of each DAC. You can choose to ignore timing faults; the sequence will run with the changes as close to the requested time as possible.

Copy section The Copy button opens a dialog where you can copy items in the current section to a range of sections. This makes it easy to create lists of similar stimuli. The left-hand side of the dialog sets the items to copy; the right-hand side sets the destination. If you clear the **Section Length** checkbox, the length of the target sections does not change, otherwise the target sections are set to the length of the source section. If you check the **Control track** box, the **Section Length** box is ignored and the length is always copied. For any item selected for copying, the corresponding item in the target section is first deleted, then the source item is copied.



Setting initial DAC and digital levels To set the initial DAC and digital levels, click on a track clear of any added items. The initial digital output level is 0 or 1; the initial DAC level is in the units set by the Sequencer tab of the Sampling configuration. Every section starts with the instructions that set these levels, so output changes caused by them will be as close to time 0 in a section as possible. The fact that these levels may not be set at exactly time zero is not considered a timing fault.

Graphical palette The palette contains 11 items you can drag and drop in the graphical window to generate a sequence. There are three dropping zones: a digital output, a DAC output and the control track. Items will only drop onto suitable targets. For example, you cannot drop a sinusoid on a digital track. There is no need to drop the items at exactly the right time point; you can edit the position afterwards.



Dragging and duplicating items You can click on a dropped item and drag it to a new target position. Hold down Ctrl as you release the item to drop a duplicate at the target position.

Common editable fields When you select an item in the graphical window you can edit the fields that relate to it. The following are described here to avoid repeating the descriptions.

Label This field is normally blank. Use it to label the selected item so that you can branch to it. A label can be up to 6 alphanumeric (A-Z, 0-9) characters long and is case insensitive; abc23 and ABC23 are the same. Labels must be unique in each section, but you can have the same label in a different section.

You cannot set a label for the initial levels of the digital and DAC tracks or for the control track as these items all start at time 0 and can be branched to by referring to the section name. You cannot set a label for arbitrary waveform output as this would prevent an important optimisation required when the sequence is generated.

Start time All digital and DAC change items have a **Start time** field and all items on the control track have an **At** field. The time is in units of seconds or milliseconds, as set by the **Graphical editor settings**, and is relative to the start time of the sequence section. There is a spin control to nudge the time on or back by the sequencer time resolution.

If you add items to the control track that take an unknown time to complete, for example a random delay or a wait for an input signal to achieve a set value, the **At** time determines where the items are drawn in the graphical editor. In this case the sequencer will maintain the time intervals between items wherever possible.

For pulses, ramps, sinusoidal and arbitrary waveform output, the sequencer attempts to produce the first output change at exactly the time you specify. For other item types, the sequencer attempts to run the first instruction of the item at the specified time.

Length Several items have a length in seconds or milliseconds. For all the pulse types, this is the period for which the pulse changes to the amplitude or level you set before it reverts to the original level. For ramps and sinusoids, this is the output length. For arbitrary waveform output, this is set to the length of the arbitrary wave you associate with the command. You can set it to be less than the length of the wave, in which case the output is stopped after the time you have set. There is a spin control to nudge the time on or back by the sequencer time resolution.

Interval For the pulse train items, this is the time between pulse starts in seconds or milliseconds.

Level (units) These fields are used with DAC pulse outputs. The field you see depends on the state of the **Absolute pulse levels** checkbox in the **Graphical editor settings**. The **Size** field sets the amplitude of a pulse; the **Level** field sets the absolute level of a pulse. You set the value in the DAC units set in the **Graphical editor settings**.

No return Normally, output from single pulses, ramps and sinusoids returns to the background level at the end of the item. If you check this box, the output change is not removed at the end of the item. To indicate this, the grey rectangle surrounding the item extends to the end of the section. You could use this to ramp a DAC up to a value and leave it there.

Overlapping items If digital pulses overlap, the result is the logical OR of the pulses. If DAC items overlap on the same channel, the output depends on the state of the **Absolute pulse levels** checkbox in the **Graphical editor settings**. If it is clear, the result is the sum of the outputs. If it is checked, the output level is set by the last item in the overlap. There is an exception: arbitrary waveform output overrides all other items.

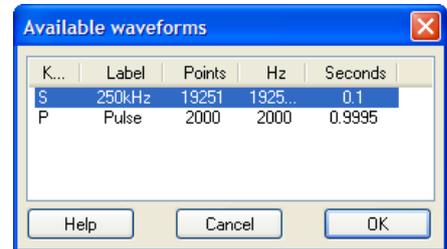
When adding pulses and pulse trains where the result would exceed the DAC range, the output is limited to the DAC range. However, pulses with an amplitude change on repeats can exceed the DAC range and wrap around. A value that goes off the top of the DAC range will reappear at the bottom; a value that goes off the bottom reappears at the top.

Single pulse  You can drag this item to either the digital or DAC outputs. For a digital output, it sets the output to be the inverse of the initial level set for this output in this section. For a DAC output, you set the amplitude of the pulse with the **Size** field or the absolute level with the **Level** field.

- Single pulse, amplitude change on repeat**  You can drag this item to a DAC output for use in a repeated section. The **Size/Level** field sets the initial amplitude or absolute level of the pulse the first time the section runs. The **Change** field sets the amplitude change to apply on each repeat. You set the number of repeats by selecting the control track and editing the **Repeats** field. The changes are calculated in real time in the 1401. If the initial level plus the number of changes times the number of repeats exceeds the DAC range, the output wraps around.
- Ramp**  You can drag this item to a DAC output. The **From** and **To** fields set the initial and final amplitudes or levels of the ramp depending on the state of the **Absolute pulse levels** checkbox in the **Graphical editor settings**.
- With a 1401*plus*, no other activity is allowed while a ramp is generated except a sinusoid. Any item that starts within the time range of a ramp will cause a timing error to be flagged and the item start will be delayed until after the ramp when you run the sequence. There is no such restriction for the other members of the 1401 family.
- Sinusoid**  You can request sinusoids on DACs 0 to 7. However only a Power1401 can use all 8 DACs. A 1401*plus* can only generate them on DACs 2 and 3, Micro1401s can only generate them on DACs 0 and 1.
- The **Size (units)** field sets the sinusoid amplitude; this is not affected by the **Absolute pulse levels** checkbox. You can offset the sinusoid with the **Centre (units)** field. If the **Absolute pulse levels** checkbox is clear, the sinusoid and offset are added to the DAC value. If the checkbox is set, the DAC output is defined by the sinusoid and offset.
- The **Period** field sets the time for one cycle of the sinusoid in seconds or in milliseconds. The **Start phase** field sets the initial phase in degrees. The output is a cosine, so a phase of 0 means start at maximum amplitude. A phase of -90 or 270 produces a sine output.
- Pulse train**  You can drag this to a DAC or digital output. It generates a train of pulses defined by the number of pulses (**Pulses** field), the length of each pulse (**Length** field) and the interval between pulse starts (**Interval** field). For a DAC output you can also set the amplitude with the **Size/Level** field.
- This method is inefficient if you need to generate a large number of pulses as each pulse takes several instructions. You should consider setting up a single pulse as a section and repeating the section to create the pulse train.
- Pulse train with varying amplitude**  You can drag this to a DAC output. It generates a train of pulses defined by the number of pulses (**Pulses** field), the length of each pulse (**Length** field) and the interval between pulse starts (**Interval** field). There is no gap after the final pulse in the train. You set the amplitude of the first pulse with the **Size/Level** field. The pulse size changes by the value in the **Change** field for each repeat.
- If you need to generate a large number of pulses it is more efficient to create a section with a single pulse with changing amplitude and repeat it to create the pulse train.
- Arbitrary waveforms**  Drag this item to the control track to start playing a waveform through the 1401 DAC outputs. You can play arbitrary waveforms through any of the DACs in your 1401. The DACs used are set when you create the waveform and can be edited in the **Play Waveform** tab of the **Sampling configuration**. Create the arbitrary outputs with the **Sample menu Output Waveform** command or from a script.
- If more than one waveform is defined and enabled in the sampling configuration, you are prompted to choose one from a list. Double-click the dropped item in the control track to

redisplay the list. Each waveform is identified by a code; this is either a single alphanumeric character or two hexadecimal digits. When you set this in the **Key** field or by selecting a waveform, the **Length** field is set to the waveform length or to a lesser value if the wave is longer than the section or to 0 if there is no matching wave. Set the length shorter than the entire wave to truncate the output.

Arbitrary output takes appreciable time to set it up; times in excess of 10 milliseconds are possible. When generating the output, this set up is moved as far forward in the sequence section as possible so that the output starts at the exact time you set. If the preparation has not completed by the time you request output to start, the sequence stalls until it is ready.



When the arbitrary output ends, the DAC outputs are returned to the background level as soon as possible. If the arbitrary waveform has more than one channel, there will be one sequencer clock period between each DAC changing to the background level.

Wait for time, condition or variable

You can drag this item to the control track only. This item pauses the sequence until a specified condition is met. Because the duration of this item may not be known, it is drawn as if it was of zero width. You can select the following delay types:

Fixed delay Set the time to wait in the **Delay** field. You will find this useful if you have short periods of sequencer activity separated by long delays.

Random delay This holds the sequence for a time between **Min** and **Max** seconds or milliseconds. All delays between the minimum and maximum have equal probability (within the capabilities of the sequencer). You will get the best result for delay ranges that are long compared to the time resolution.

Poisson delay This generates a delay with a Poisson statistic; the delay has the same probability of ending at any time during it. The **Time const** field sets the average delay length.

Digital input high/low You can wait for a nominated digital input bit in the range 0-7 to be high or low. These are not the same bits as used for the event inputs. If you want to synchronise to a bit changing you must wait for it to be in the opposite state to the one you want first. If you need the sequence to perform actions while you wait, use a branch item. You can wait for combinations of input bits with the text sequencer.

Channel above/below/outside/within You can wait for a nominated waveform channel sampled by the 1401 to be above or below a level set by the **Threshold** field, or to be outside or within a pair of levels set by the **Lower** and **Upper** levels fields. The levels are set in the channel units, as stored in the sampling configuration. If you subsequently change the channel type, the results will not be harmful, but the sequence will not operate as intended!

Next event You can wait for the number of events set by the **Count** field to occur on a nominated event, marker or **WaveMark** channel that is sampled by the 1401. When this delay item is reached, the sequence notes how many events have been sampled on the channel and waits until the number increases by the count.

Time reached This item waits until the sample time reaches the value you set. If the sample time has already passed this time, there is no wait.

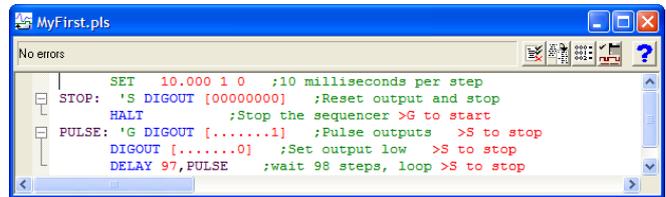
Cosine phase 0 This item waits for the next time that cosine output on the nominated DAC channel passes through phase zero.

- Variable comparisons* You can wait for various conditions based on sequence variables. You can use variables 1 to 25 of the sequencer variables. The remaining variables are used to implement the sequence sections. Variable values are 32-bit integers. You can manipulate the variables in the **Variable arithmetic** item. Variables are described in detail for the text sequencer.
- Event burst* This item monitors an event, marker or WaveMark channel sampled by the 1401 for a group of events with a user-defined maximum separation. The **Intervals** field sets the number of gaps to check and the **Max time** field sets the maximum acceptable interval. If any interval is greater than the maximum, the sequence starts the search again. For this to work well, the maximum interval must be significantly greater than the time resolution of the sequence.
- Branch on condition, probability or variable**  With this item you can break the normal flow of the sequence and branch to a different section or to a label you have defined for an item in the current section. All branches have a **Branch destination** field in which you can select a section to branch to, or you can type the name of a label in the current section.
- When you branch, the timing to the target may not be exactly what you expect. The sequence will take one or more steps to implement the branch and the target instruction may require preparatory steps. Such effects are small unless you use arbitrary waveform output where the preparatory steps can take several milliseconds. If you need the tightest possible control over branch timing you should consider using the text sequence editor. The branches you can set are:
- Probability* **Percent** sets the probability of the branch, 0% never branches, 100% always branches.
- Digital input high/low* The **Bit** field sets the digital input bit number in the range 0-7 to test.
- Channel above/below/outside/within* You can branch if a nominated waveform channel sampled by the 1401 is above or below a level set by the **Threshold** field, or is outside or within a pair of levels set by the **Lower** and **Upper** levels fields. The levels are set in the channel units, as stored in the sampling configuration. If you subsequently change the channel type, the result is not harmful, but the sequence will not operate as intended!
- Variable comparisons* You can branch on the result of comparing sequence variables with constant values and other variables. Variable values are 32-bit integers. You can manipulate the variables in the **Variable arithmetic** item. Variables are described in detail for the text sequencer.
- Unconditional* This always branches to the destination.
- Time comparisons* You can compare a variable plus a time offset with the current time and branch on the result. You can set the variable to the current time (plus a time offset) with the variable arithmetic *current time* instruction.
- Response with timeout* You can wait for a new data item in an event, marker or WaveMark channel sampled by the 1401. The branch is taken if a new item is detected within the timeout period.
- Generate digital marker channel event**  This adds an event to the Marker channel (if it is enabled). You can set the marker code with the **Code** field or check the **Record data** box to record the state of digital input bits 0 to 7 (these are not the same bits used for event inputs). The **Code** field should be set to one character or to two hexadecimal digits.
- Variable arithmetic**  Although the use of variables is more commonly done with the text editor, you can perform basic variable manipulation here. You can use variables 1 to 25 for these operations. Variable values are 32-bit integers. In all cases, the variable that is changed is set by the **Target var** field. Operations are:

- Set to value/variable* Replace the target variable with the contents of the Value field or of a variable.
- Add/subtract value/variable* Adds or subtracts the contents of the variable or value.
- Multiply by value/variable* Multiplies the target variable by the variable or value.
- Random value* This replaces the target variable by a random number that is from 1 to 30 bits long, set by the Bits field. The possible values for an n bit number are 0 up to 2ⁿ minus 1. For example, if the Bits field is 4, the possible results are 0 to 15.
- Current time* The variable is set to the current sample time plus a time offset, in Spike2 clock ticks, as set in the Resolution tab of the Sampling configuration.
- Fixed time* The variable is set to a time, in Spike2 clock ticks, as set in the Resolution tab of the Sampling configuration.

The text editor

The text output sequence is stored as a text file with the extension .PLS. You can open existing Spike2 output sequence files or create new ones with File menu New... (see the File menu chapter for details of file types). This picture shows an output sequence that has been typed in without formatting. There are five buttons at the top of the window:

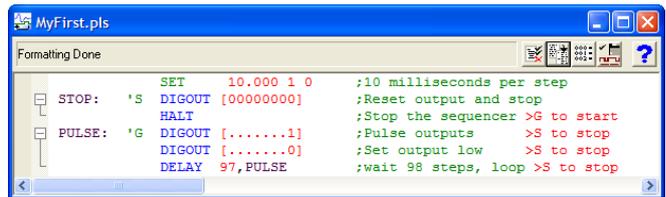


This picture shows an output sequence that has been typed in without formatting. There are five buttons at the top of the window:

Format



This aligns the labels, key codes, instructions and any arguments, output text and comments and removes step numbers. Undefined labels are not flagged but there must be no other errors. The picture shows the result.

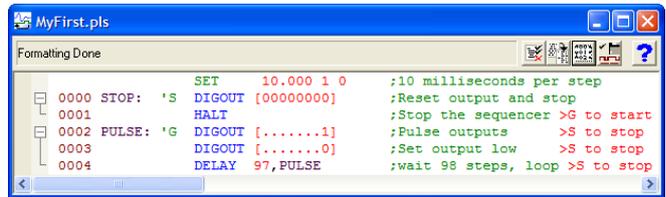


The picture shows the result.

Format with step numbers



This does the same job as the Format button, and also starts each line with the step number. Step numbers can be useful as they give an indication when you are running out of space and can pinpoint the line where your sequence is not behaving as you expect.

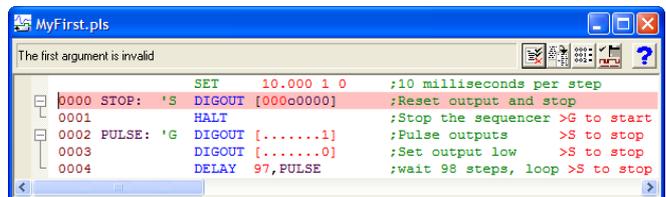


can pinpoint the line where your sequence is not behaving as you expect.

Compile



This checks the sequence to make sure that it is correct with no labels missing or duplicated and no duplicated key codes. The picture shows a sequence with a simple error (the 0 in the second line should be a zero). The line in error is marked and an explanatory message is shown at the top of the window.



Current  This compiles the sequence and if it compiles without error, the sequence is saved. If there is any error, it is reported and the command ends. If there is no error, what happens next depends on whether Spike is sampling data or not:

Not sampling The name of the sequence file is set in the sampling configuration and the sequence will be used when you next sample data. You can also set the output sequence file from the **Sampling Configuration** dialog accessed from the **Sample** menu.

Sampling Spike attempts to replace the current sequence. This will fail if there is no current sequence, or if the sequence or table size is larger than the current size or the size reserved in the **Sequencer** tab of the **Sampling configuration** dialog. If a sequence is replaced, it runs from the first instruction. Any variables that are declared with initial values are updated as are all initialised table values.

A message indicating the outcome of the operation is displayed at the top of the window.

Help  This is the Help button. It opens a window holding a list of the sequencer topics for which help is available. You can copy and paste text from the help window into your sequence. You can also look up the word under the text caret with the F1 key.

Loading sequencer files for sampling The name of the output sequencer file to use during sampling is part of the sampling configuration. The file name, including the path to the folder containing it, must be less than 200 characters long. You set the file name either with the **Current** button, as described above, or in the **Sampling Configuration** dialog. When you start sampling, Spike2 searches for any output sequence file named in the sampling configuration. Spike2 compiles output sequence files whenever you use them. If a file contains errors you are warned and the file is ignored. If the source file is write protected or on a read only medium, you are not allowed to edit the sequence in Spike2.

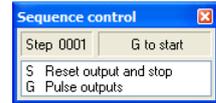
Getting started The sequencer runs instructions in order unless told to branch. It can be re-routed during data acquisition by associating an instruction with a key on the keyboard. Each time the key is pressed or the associated sequencer control panel button is clicked or the script language `SampleKey()` command is used, the sequencer jumps to the associated instruction. Spike2 records keys pressed during sampling, so you have a record of where in the document you switched to a new portion of the sequence.

Here is a simple sequence that will pulse digital output bit 0 for 10 milliseconds once per second. You can start and stop the pulses with keys or by clicking buttons. You will find this in `Spike6\Sequence\MyFirst.pls` to save you typing it in.

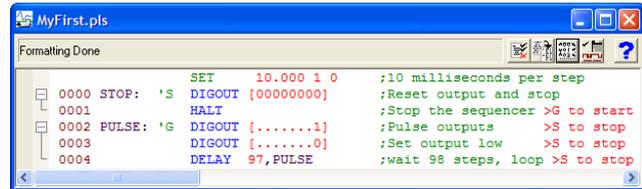
```
          SET      10          ;10 milliseconds per step
          'S DIGOUT [00000000] ;Reset output and Stop
          HALT          ;Stop the sequencer >G to start
PULSE:   'G DIGOUT [...1]   ;Pulse outputs >S to stop
          DIGOUT [...0]    ;Set output low >S to stop
          DELAY 97,PULSE   ;wait 98 steps, loop >S to stop
```

Open `MyFirst.pls` and click the **Check and make current sequence** button at the upper right of the window (leave the mouse over each button for a second or so to see the descriptive text). Open the **Sampling configuration** dialog, and set a configuration with one event channel on port 0. To record the output pulses you must connect the digital output to the event input. This is easy with the **Micro1401** and **Power1401** as you can connect Digital output 0 to Event input 0 on the front panel. For the **1401plus** you need to connect the digital output pin 17 to the digital input pin 17. You do not need to make the connection to follow this description.

Click the Run now button in the sampling configuration, and then click the Start button. The sequencer control panel is now visible. It displays Step 0001 in the top left window and G to start to the right. If the control panel is docked there are two buttons labelled S and G, otherwise you can see the text S Reset output and Stop and G Pulse outputs. To start the output, click on the G or type G on the keyboard. While it runs, the display will change to Step 0004 (with occasional flickers) and S to stop. If you have connected the digital output to the input, you will see your pulses recorded, once per second. Click the S or type s on the keyboard to stop the output.



As this is our first sequence we will explain it in detail. Click the Format and add step numbers button at the top of the sequencer window. All the lines get a step number except the first.



This is the step number displayed in the control panel. You can remove step numbers with the Format button. Lines that get a number are part of the sequence run by the 1401. The SET line tells the sequencer how fast to run, in this case 10 milliseconds per step, and is not part of the sequence run by the 1401. A SET or SCLK directive is usually the first line in the sequence.

The next line (step 0000) is the first to run when you start sampling. The 'S means *jump to this line every time the S key is pressed during sampling* (as long as the data window is the current window). To allow you a wide choice of keyboard jumps, lower case s and upper case S are treated as different keys.

The instruction on this line is DIGOUT [00000000], which sets 8 digital outputs (bits 15 to 8 of the digital output) to the low state, nominally 0 Volts. The remainder of the line is a comment. Because there is a keyboard jump set for this line, the comment is also used as a label for the sequencer control panel.

The next line, (step 0001) holds a HALT instruction. This stops the sequencer and nothing will happen until the sequencer is told to jump to one of the steps labelled with a key code. The sequencer control panel displays the text to the right of the > for the current step. When the sequencer halts, it stays at step 0001, so you see the message G to start.

Steps 0002 and 0003 set digital output bit 8 high (nominally 5 Volts) and low again. The sequencer is running at 10 milliseconds per step, so the pulse is 10 milliseconds wide. Step 0002 starts with a label, PULSE:, so we can branch back here in the next step.

Step 0004 has the instruction DELAY 97,PULSE to make the sequence wait for 97 extra step times in addition to the step time that every instruction takes, and then branch to the instruction labelled PULSE. This instruction takes 980 milliseconds, and together with the 20 milliseconds taken by steps 0002 and 0003, this makes 1000 milliseconds for the loop. Instead of 97 we could have written s(1)-3. The s(1) function returns the number of sequencer steps in 1 second.

The remainder of this chapter is organised as a reference manual for the sequencer instructions. You can find more information about the output sequencer in the *Spike2 Training Course* manual.

Instructions		
		These are the output sequencer instructions in Spike2 versions 5, 6 and 7. Instructions in brackets are obsolete and should be avoided in new sequences.
<i>Digital (TTL compatible) input and output (see page 5-4)</i>	DIGOUT DIGLOW DIBNE,DIBEQ DISBNE,DISBEQ WAIT (DIGIN) (BZERO,BNZERO)	Write to digital output bits 15-8 Write to digital output bits 7-0 (does nothing in 1401 <i>plus</i>) Read digital input bits 7-0, copy to V56 test and branch Test last read digital inputs (in V56) and branch Wait for a pattern on the lower 8 digital inputs Read and test digital inputs Branch as result of DIGIN test
<i>DAC (waveform/voltage) outputs (see page 5-27)</i>	DAC, DACn ADDAC, ADDACn RAMP	Set DAC value for DACs 0-7, DACn for DACs 0-3 Increment DAC by a value Set automatic DAC ramping to a target value
<i>Sinusoidal waveform output (see page 5-29)</i>	SZ SZINC RATE RATEW ANGLE WAITC RINC RINCW PHASE OFFSET CLRC	Set the cosine output amplitude (CSZ, DSZ) Change the cosine output amplitude (CSZINC, DSZINC) Set the cosine angular increment per step (CRATE, DRATE) As RATE, but waits for phase 0 (CRATEW, DRATEW) Set cosine angle for the next step (CANGLE, DANGLE) Branch until cosine phase 0 (CWAIT, DWAIT) Change the cosine angle increment per step (CRINC, DRINC) As RINC but waits for phase 0 (CRINCW, DRINCW) Defines what phase 0 means (CPHASE, DPHASE) Offset for sinusoidal output (COFF, DOFF) Clear the new cycle flag (CWCLR, DWCLR)
<i>General control (see page 5-35)</i>	DELAY DBNZ (LDCNTn, DBNZn) Bxx CALL CALLV, (CALLn) RETURN JUMP HALT NOP	Do nothing for a set number of steps Decrement a variable and branch if not zero Load counter 1 to 4 (V61-V64), decrement, branch if not zero Compare variables and branch (xx = GT, GE, EQ, LE, LT, NE) Branch to a label, save return position Like CALL, but load a variable (counter 1-4) with a value Branch to instruction after last CALL, CALLV or CALLn Unconditional branch to a label Stops the sequencer and waits to be re-routed This does nothing for one step (NO OPERATION)
<i>Variable arithmetic (see page 5-37)</i>	ADD, SUB ADDI, (ADDIL) MOV MOVI, (MOVIL) MUL,MULI NEG, ABS DIV, RECIP	Addition and subtraction of variables Add a constant value to a variable Copy one variable to another Move a constant value into a variable Multiply two variables, multiply by a constant Negate variable, absolute value of variable Division and reciprocal of variables
<i>Variable logic (see page 5-40)</i>	AND, ANDI OR, ORI XOR, XORI	Bitwise AND of variables, variable and constant Bitwise OR of variables, variable and constant Bitwise exclusive OR of variables, variable and constant
<i>Table support (see page 5-40)</i>	TABLD, TABST TADADD, TABSUB TABINC	Load a register from the table and store a register to the table Add a table value to a variable or subtract it from a variable Increment a register and branch while within the table
<i>Access to data capture (see page 5-42)</i>	REPORT, MARK CHAN TICKS	Simulate an external E1 pulse to record/set a digital marker Get the latest waveform value or event count from a channel Load a variable with the time in Spike2 time units
<i>Randomisation (see page 5-43)</i>	BRAND, (BRANDV) MOVRND (LD1RAN)	Random branch with a probability Load a variable with a random number Load counter 1 (V61) with a random number 1-256

<i>Arbitrary waveform output</i> (see page 5-44)	WAVEGO	Start or prepare arbitrary waveform output area
	WAVEBR	Test arbitrary waveform output and branch on the result
	WAVEST	Start or stop arbitrary waveform output

Instruction format

Blank text lines and lines with a semicolon as the first non-blank character, are ignored. Instructions are not case sensitive. Each instruction has the format:

```
num lab:      'key  code  arg1,arg2,... ;Comment      >display
```

num An optional step number in the range 0 to 8190, for information only.

lab: An optional label, up to 8 characters long followed by a colon. The first character must be alphabetic (A-Z). Labels are not case sensitive. Labels may not be the same as instruction codes or variable names.

'key In this optional field, *key* is one alphanumeric (a-z, A-Z, 0-9) character. When this character is recorded as a keyboard marker during data capture, the sequencer jumps to this instruction. Each *key* can occur once. Upper and lower case are distinct. The *key* appears in the sequencer control panel.

code This field defines the instruction to be executed. It is not case sensitive.

arg1,... Instructions need up to 4 arguments and are separated by commas or spaces. These are described with the instructions. If an argument can be represented in different ways, they are separated by vertical bars (read as "or"), for example: *expr|Vn|[Vn+off]*. In this case, the argument can be an expression, a variable or a table reference.

comment The text after the semicolon is to remind you of the reason for the instruction. If a *key* is set, this comment also appears in the sequencer control panel.

>display When a sequence runs, text following a > in a comment is displayed in the sequencer control panel to indicate the current instruction. The special combination >" means the same display as the previous instruction and the combination >= means no change to whatever is currently displayed.

Expressions

Many instructions allow the use of an expression in place of a constant value, indicated by *expr*. An expression is formed from constants, numbers, labels, round brackets (and), the operators +, -, * and /, and sequencer expression functions. Labels have the value of the instruction number in the sequence; the first instruction is numbered 0.

The operators * and / (multiply and divide) have higher priority than + and - (add and subtract). This means that $1+2*3$ is interpreted as $1+(2*3)$ and not as $(1+2)*3$. Apart from this, evaluation is from left to right unless modified by brackets.

The sequence compiler evaluates expressions as real numbers, so $3/2$ has the value 1.5. If *expr* is used as an integer, for example `DELAY expr`, it is rounded to the nearest integer. Values in the range 3.5 to 4.49999... are treated as 4. Before version 4.06 the result was truncated, so 3.0 to 3.9999... was treated as 3.

Sequencer expression functions

These functions can be used as part of expressions to give you access to Spike2 clock and sequencer step timing and to convert between user units and DAC and ADC values.

`s(expr)` The number of sequencer steps in *expr* seconds, milliseconds and
`ms(expr)` microseconds. For example, with a step size of 200 milliseconds,
`us(expr)` `s(1.1)` returns 5.5. This is often used with the `DELAY` instruction. Each
instruction uses 1 step, so use `DELAY s(1)-1` for a delay of 1 second.

`sTick(expr)` The number of Spike2 sample clock ticks in *expr* seconds,
`msTick(expr)` milliseconds and microseconds. The result is in the same units as

usTick (expr)	returned by the TICKS instruction, so TICKS V1, sTick(1) sets V1 to the value that TICKS will return 1 second later.
Hz (expr)	The angle change in degrees per step for a cosine output of expr Hz. For a 2 Hz cosine, use RATE n, Hz (2). To slow by 0.1 degrees per step use RINC n, Hz (-0.1). Use in RATE, RATEW, RINC and RINCW.
VHz (expr)	The same as Hz (), but the result is scaled into the 32-bit integer units used when a variable sets the rate. MOVI V1, VHz (2) followed by CRATE V1 will set a 2 Hz rate.
VAngle (expr)	Converts an angle in degrees into the internal angle format. The 32-bit integer range is 360 degrees. The result is expr * 11930464.71.
VDAC16 (expr)	Converts expr user DAC units so that the full DAC range spans the full range of a 16-bit integer. Use with variables for DACn and ADDACn.
VDAC32 (expr)	Converts expr user DAC units into a 32-bit integer value such that the full DAC range spans the 32-bit integer range. Use this to load variables for use with the DAC and ADDAC instructions.
ASz (expr)	ASz converts user DAC units into values for use with SZ and SZINC.
VSz (expr)	VSz converts user DAC units into variable values for SZ and SZINC.
TabPos ()	The number of table data items defined at this position by TABDAT.
DRange ()	The DAC output range in Volts, 5.0 for ±5 or 10.0 for ±10 Volts.

Used with care, the built-in functions allow you to write sequences that operate in the same way regardless of the sequencer step time or DAC scaling values.

Variables

You can use the variables, V1 to V256 (V64 is the maximum for a 1401*plus*), in place of fixed values in many instructions. In the sequencer command descriptions, Vn indicates the use of a variable. Where a variable is an alternative to a fixed value expression we use expr|Vn. Variables hold 32-bit integer numbers that you can set and read with the SampleSeqVar () script command.

Some variables have specific uses: Variables V57 through V64 hold the last value written by the sequencer to DACs 0 through 7, V56 holds the last bit pattern read from the digital inputs with the DIBxx or DIGIN instructions, V61 through V64 are also used to emulate the counters for the LDCNDn and DBNZn instructions used in previous versions of Spike2.

VAR directive

You can assign each variable a name and an initial value with the VAR directive. Names must be assigned before they are used, usually at the start of the sequence. The syntax is:

```
VAR    Vn, name=expr           ;comment
```

VAR does not generate any instructions. It makes the symbol name equivalent to variable Vn and sets the initial value when the sequence is loaded. Anywhere in the remainder of the sequence where Vn is acceptable, name can be used. name can be up to 8 characters, must start with an alphabetic character and can contain alphabetic characters and the digits 0 to 9. Variable names are not case sensitive. A variable name must not be the same as an instruction code, a label, a built-in function name or a constant.

There is no need to specify a name or an initial value. If no initial value is set, a variable is initialised to 0 even if not included in a VAR statement. Spike2 automatically assigns V56 the name VDigIn and variables V57 through V64 the names VDACC0 through VDACC7. The following are all acceptable examples:

```
VAR    V1, Wait1=ms(100)      ;Set name and initial value
VAR    V2, UseMe              ;Set name only, so value is 0
VAR    V3=200                 ;No name, initialise to a value
VAR    V4                     ;No name, initialised to 0
```

When a variable is used in place of a bit pattern in a digital input or output instruction, bits 15 to 8 and bits 7 to 0 have different uses. In the expressions that describe these operations we write $V_n(7-0)$ and $V_n(15-8)$ to describe which bits are used. BAND means bitwise binary AND (if both bits are 1, the output is 1, otherwise 0), BXOR means bitwise exclusive OR (if both bits are different the output is 1, otherwise 0).

When used in one of the cosine output angle instructions, the 32-bit variable range from -2147483648 to 2147483647 represents -180° up to $+180^\circ$. The VarValue script in the Scripts folder calculates variable values for the digital and the cosine instructions.

Script access to variables Scripts can set and read the variable values with the SampleSeqVar() script command. See *The Spike2 script language* manual for details. You can set initial values from the script as long as you set the values after you create the new data file, but before you start sampling. Values set in this way take precedence over values set by the VAR directive.

Constants

It is often useful to define a numeric value as a named constant. For example, when referencing a table value $[V1+Slope]$ is easier to understand than $[V1+23]$. It also helps to maintain code; if you need to change a numeric value that is used often make it a named constant and just change the constant once. The = directive does not generate any instructions, it just makes a name equivalent to a number, and the name can be used anywhere a numeric expression is valid. The syntax is:

```
name = expr ;comment
```

The name must start with an alphabetic character and can contain alphabetic characters and the digits 0 to 9. The name can be up to 8 characters long. Numeric constant names are not case sensitive and must not clash with label, instruction, variable or built-in function names. Examples of use:

```
Wait1 = ms(100)
Wait2 = Wait1*1.3
oSlope = 23
DELAY Wait1-1 ;Use constant in an instruction
ADDTAB V1,[V2+oSlope] ;Use as a table offset
```

The expression values are calculated and stored as floating point numbers. If they are used in a context that requires an integer, the fractional part of the number is ignored.

The = directive was added at version 7.00.

The SET, SCLK and SDAC directives

These directives control the interval between sequencer clock ticks and the DAC output units. Directives are not part of the sequence and do not occupy a step. These directives should occur before any step-generating code. The SCLK and SDAC directives were added at version 6.03 to separate setting the tick rate from setting the DAC output units.

```
SET msPerStep, DACscale, DACoffset
SCLK msPerStep
SDAC DACscale, DACoffset
```

The msPerStep field sets the milliseconds per step in the range 0.010 to 3000. The table shows the minimum step times and timing resolution for each type of 1401. To ensure accurate timing, msPerStep must be an integral multiple of Resolution or Spike2 will not sample. This is checked each time Spike2 samples data.

	Power	Micro1401	micro1401	1401plus
Minimum step (ms)	.010	.010	.050	.050
Resolution (ms)	.001	.001	.004, .006, .010	.004, .006, .010

If there is no `SET` or `SCLK` directive, the sequence runs at the default rate, which is 10 milliseconds per step unless the script command `SampleSeqClock()` has changed it. If there is no `SET` or `SDAC` directive, `DACScale` is 1 and `DACOffset` is 0. The sequencer expression functions `s()`, `ms()`, `us()`, `Hz()` and `VHz()` use the `msPerStep` value. If you use these functions, the `SET` or `SCLK` directive must occur first in the sequence.

DAC scaling The DACs in the 1401 are implemented so that the full range maps onto the full range of 16-bit signed integer numbers (-32768 to +32767). If you have 16-bit DACs a change of 1 changes the output by the smallest step possible. With 12-bit DACs, the smallest step is a change of 16. For most purposes, it is easier to work in units such as Volts or millivolts rather than these DAC units. However, the 1401 works in DAC units and if you set DAC values using variables, the variable values are based on DAC units. In the `SET` and `SDAC` directives, the `DACscale` and `DACoffset` fields define the conversion from user units into DAC units. The standard values of 1.0 for the scale and 0.0 for the offset make the full scale DAC values run from -5 to +4.99985. As most systems have ± 5 Volt analogue systems, the standard scale and offset let you work with the DACs in Volts.

`DACScale` The number of user units that correspond to an output change of 6553.6 DAC units (1 Volt for a ± 5 Volt system, 2 Volts for a ± 10 Volt system). The standard value 1.0 is used if you omit `DACScale`.

`DACOffset` The user units that correspond to 0 DAC units (0 Volt output). The standard value 0.0 is used if you omit `DACOffset`.

Please remember that `DACScale` and `DACOffset` do not change the DAC outputs in any way. They are a convenience to allow you to enter values in units that are appropriate to your application. The sequencer expression functions `VDAC16()` and `VDAC32()` use `DACScale` and `DACOffset`, so they must come after the `SET` or `SDAC` directives.

For example, consider the case where the DACs drive a patch clamp amplifier where a change of 2.5 Volts into the amplifier causes a 200 mV potential at the cell and 0 Volts into the amplifier is 0 mV at the cell. For a ± 5 Volt system, 2.5 Volts is 16384 DAC units, so `DACScale` is $200 * 6553.6 / 16384$, which is 80. `DACOffset` is 0, as an output of 0 produces 0 mV at the cell.

Table of values From Spike2 version 5 onwards the sequencer supports a table of 32-bit values. You declare a table with the `TABSZ` directive, normally at the start of your sequence:

```
TABSZ expr
```

Where `expr` is an expression that sets the number of items in the table. The table size must evaluate to a number in the range 1 to 1000000. Each table item is a 32-bit integer and uses 4 bytes of 1401 memory. The maximum size in a 1401 with 1MB of memory, and assuming that there is no arbitrary waveform output, is around 150000 items. The first table item has an index number of 0, the second item is index 1, and so on.

Setting table data From the script language you can move data between an integer array and the table with the `SampleSeqTable()` function. You can also preset table data from the sequence with the `TABDAT` directive, which must come after the `TABSZ` directive:

```
TABDAT expr
TABDAT expr,expr,expr...
```

Where `expr` is an expression that evaluates to a 32-bit integer. Each `TABDAT` directive adds data to the table, starting at the beginning. The sequencer compiler flags an error if you define more data than set by `TABSZ`. Table data declared in this way is transferred to the 1401 when you create a new data file to sample. If you do not set the table data with the `TABDAT` directive or from a script, the table values are undefined. The `TabPos()`

expression has the value of the number of data items that have been defined at the point where it is used. For example:

```
TABDAT TabPos() ; a table item that holds its own index.
```

Accessing table data Although you can move data between one of the variables and the table with the `TABLD` and `TABST` instructions, many instructions access the table directly. It takes more time to use a table than to use a variable.

All references to the table use the contents of one of the variables as an index into the table plus an optional offset as: `[Vn]` or `[Vn+off]` or `[Vn-off]`. The offset `off` is an expression that evaluates to a number in the range `-1000000` to `1000000`. For example, if `V1` holds 10, `[V1]` refers to the contents of index 10, `[V1-10]` refers to index 0 and `[V1+10]` refers to index 20. Out of range table indices read 0 and are non-destructive.

The `TABINC` instruction makes it easy to increment a variable used as a table index and branch until the increment generates an index outside the table. The following example generates five DAC outputs at 5 different intervals. The final instruction uses `oNext` as the offset from one pair of entries to the next and `oDelay` and `oDac` as offsets to the delay and DAC values. In this trivial case we could just as easily used the values 2, 0 and 1. However, if the table encoded a repeating list of 10 values, using named constants in place of numbers would be much clearer.

```

                SET      1,1,0          ; lms per step, normal scales
oDelay          TABSZ   10              ; table of 10 items
                =       TabPos()       ; offset to delay
                TABDAT  ms(1000)-3     ; Delay value (1000 ms)
oDac            =       TabPos()       ; offset to DAC value
                TABDAT  VDac32(1)      ; DAC value (1 Volt)
oNext           =       TabPos()       ; increment to next item
                TABDAT  ms(100)-3,VDac32(2) ; 100 ms, 2 Volts
                TABDAT  ms(50)-3,VDac32(3) ; 50 ms 3 Volts
                TABDAT  ms(500)-3,VDac32(-1) ; 500 ms -1 Volt
                TABDAT  ms(200)-3,VDac32(0) ; 200 ms 0 Volts

TLOOP:          MOVI   V1,0             ; use V1 as table index, set 0
                DELAY  [V1+oDelay]     ; programmed delay
                DAC    0,[V1+oDac]     ; set DAC 0 to the value
                TABINC V1,oNext,TLOOP; add 2 to V1, branch if in table

```

Long data sequences To output a very long data sequence, you can use the table as a buffer. The basic idea is to divide the table into two halves and use a script to transfer new data into the half of the table that the sequence is not using. To find out where the sequence has reached, look at the value of the variable used as an index with `SampleSeqVar()`. Set a large enough table size so that the time taken to use half the table is several seconds.

Include files There are times when you will want to reuse definitions or sequence code sections in multiple projects. You can do this by pasting the text into your sequence, but it can be more convenient to use the `#include` command to include files into a sequence. A file that is included can also include further files. We call these *nested* include files. Only the first `#include` of a file has any effect. Subsequent `#include` commands that refer to the same file are ignored. This prevents problems with output sequence files that include each other and stops multiple definitions when two files include a common file. A `#include` command must be the first non-white space item on a line. There are two forms of the command:

```
#include "filename" ;optional comment
#include <filename> ;optional comment
```

where `filename` is either an absolute path name (starting with a `\` or `/` or containing a `:`), for example `C:\Sequences\MyInclude.pls`, or is a relative path name, for example `include.pls`. The difference between the two command forms lies in how relative path names are treated. The search order for the first form starts at item 1 in the following list. The search for the second form starts at item 3.

1. Search the folder where the file with the `#include` command lives. If this fails...
2. Search the folder of the file that included that file until we reach the top of the list of nested include files. If this fails...
3. Search any `\include` folder in the folder in which Spike2 is installed. If this fails...
4. Search the current folder.

Included files are always read from disk, even if they are already open. If you have set the **Edit** menu **Preferences** option to save modified scripts and sequences before running, modified include files are automatically saved when you compile. If this option is not set, the output sequence compiler will stop the compilation with an error if it finds a modified include file. You must save the included file to compile your sequence.

There are no restrictions on what can be in an included file. However, they will normally contain constant and variable definitions and possibly user-defined code. It is usually a good idea to have all your `#include` commands at the start of a sequence file so that anyone reading the source is aware of the scope of the sequence.

The `#include` command for output sequences was added to Spike2 at version 7.00 and is not recognised by any version before this. A typical file using `#include` might start with:

```
#include <sysinc.pls>           ;my system specific includes
#include "include/proginc.pls" ;search script relative folder
set 1,1,0                      ;start of my code...
```

Opening included files If you right-click on a line that holds an include command, and Spike2 can locate the included file, the context menu will hold an entry to open the file. The search for the file follows that described above, except that it omits step 2.

Errors in included files If an error is detected during the compilation of an included file, an error message is displayed at the top of the original window indicating which included file has a problem, and the included file is opened (if it can be found) and the offending line is highlighted.

Changing the sequence during sampling

As long as sampling starts with a text or graphical sequence running, it can be replaced during sampling. The **Sampling** menu **Change Output Sequence** command lets you choose a text sequence file on disk and the **Current** button in the Sequence editor sets the current document as the sequence. When sampling starts, Spike2 allocates sufficient space in the 1401 for the first sequence and for any table space it uses. If you want to load sequences during sampling that are larger or needs more table space you must reserve space using the **Sequencer** tab of the **Sampling configuration** dialog.

When a new sequence loads, all DAC values, digital output values, variable and table values are preserved except variables that are given initial values by the `VAR` directive and table values that are given values with the `TABDAT` directive.

Sequencer instruction reference

Each instruction below is followed by an example. The examples show the preferred instruction format, however the system is flexible. For example, a comma should separate arguments, but a space is also accepted. The patterns used for digital ports should be enclosed by square brackets, however you may omit the brackets if you wish.

Many of these instructions allow you to use a variable or a table entry in place of an argument. In this case, the alternatives are separate by a vertical bar, for example:

```
DELAY    expr|Vn| [Vn+off],OptLB
```

This means that the first argument can be an expression, a variable or a table entry. There is no explicit documentation for the use of the table, except in `TABLD` and `TABST`. Where table use is allowed it is written as `[Vn+off]`. If you use a table value in an instruction, the effect is exactly the same as using a variable with the same value as the table entry.

Digital I/O

These instructions give you control over the digital output bits and allow you to read and test the state of digital input bits 7-0.

DIGOUT

The `DIGOUT` instruction changes the state of digital output bits 15-8 (see page 5-4 for the connections). The output changes occur at the next tick of the output sequencer clock.

```
DIGOUT  [pattern]|Vn| [Vn±off],OptLab
```

pattern This determines the new output state. You can set, reset or invert each output bit, or leave a bit in the previous state. The pattern is 8 characters long, one for each bit, with bit 15 at the left and bit 8 at the right. The characters can be “0”, “1”, “i” or “.” standing for *clear*, *set*, *invert* or *leave alone*. You may omit the square brackets, however the `Format` command will insert them.

```
DIGOUT  [...001i] ;clear bits 3 and 2, set 1, invert 0
DIGOUT  [.....i] ;invert 0 again to produce a pulse
DIGOUT  V10      ;use variable V10 to set the pattern
```

Vn With a variable the new output is: (old output `BAND Vn(7-0)`) `BXOR Vn(15-8)`. The variable equivalent of `[...001i]` is `241+256*3`, and of `[.....i]` is `255+256*1`. If you use a table value, set the same value in the table that you would use for a variable. You can use the `VarValue` script in the `Scripts` folder to calculate variable or table values.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example produces ten 1 millisecond pulses 100 milliseconds apart.

```
SET      1          ;run at 1 ms per step
MOVI     V1,10      ;V1 holds the number of pulses
LOOP:    DIGOUT     [...1] ;bit 0 high      >Pulsing
         DIGOUT     [...0] ;bit 0 low       >Pulsing
         DELAY      ms(100)-4 ;4 inst in the loop >Pulsing
         DBNZ      V1,LOOP    ;count down     >Pulsing
         HALT      ;finished          >Done
```

DIGLOW

The `DIGLOW` instruction changes the state of digital output bits 7-0 of the `Power1401` and `Micro1401` (see page 5-4 for the connections). It has no effect on a `1401plus`. Unlike `DIGOUT`, the output changes occur immediately, they do not wait for the next sequencer clock tick. You can take advantage of this to change all 16 digital outputs almost simultaneously (within a few microseconds) by using `DIGOUT` followed by `DIGLOW`.

```
DIGLOW  [pattern]|Vn| [Vn+off],OptLB
```

pattern This determines the new output state. The pattern is 8 characters long, one for each bit, with bit 7 at the left and bit 0 at the right. The characters can be “0”, “1”, “i” or “.” standing for *clear*, *set*, *invert* or *leave alone*. You may omit the square brackets, however the **Format** command will insert them.

```
DIGLOW [...001i] ;clear bits 3 and 2, set 1, invert 0
DIGLOW [.....i] ;invert 0 again to produce a pulse
DIGLOW V10      ;use variable V10 to set the pattern
```

Vn With a variable the new output is: (old output BAND Vn(7-0)) BXOR Vn(15-8). The variable equivalent of [...001i] is 241+256*3, and of [.....i] is 255+256*1. If you use a table value, set the same value in the table that you would use for a variable. You can use the **VarValue** script in the **Scripts** folder to calculate variable or table values.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

DIBEQ, DIBNE These instructions test digital input bits 7-0 against a pattern (see page 5-4 for the connections). **DIBEQ** branches on a match. **DIBNE** branches if it does not match. Both instructions copy bits 7-0 of the digital input to V56 (VDigIn), for use by **DISBEQ** and **DISBNE**. If you use the template-matched signal in the *1401plus*, input bits 7-0 may be programmed as outputs and cannot be used as inputs.

```
DIBNE [pattern]|Vn|[Vn+off],LB
DIBEQ [pattern]|Vn|[Vn+off],LB
```

pattern This is 8 characters, one for each input bit. The characters can be “0”, “1” and “.” meaning match 0 (TTL low), match 1 (TTL high) or match anything. The bit order in the pattern is [76543210]. You may omit the square brackets, however the **Format** command inserts them.

Vn With a variable the result is: (input BAND Vn(7-0)) BXOR Vn(15-8). A result of 0 is a match, not zero is not a match.

LB The destination of the branch if the input was a match (**DIBEQ**) or not a match (**DIBNE**). This label must exist in the sequence.

This example waits for a pulse sequence in which the falling edges of two consecutive pulses are less than 2*V1+2 sequencer clock ticks apart. It waits for a falling edge, waits for a rising edge with a timeout and then waits for the next falling edge with a timeout. If timed out, we start again. If the input signal has high states less than three ticks wide, or low states less than 2 ticks wide, this example may miss them.

```
WHI:    DIBNE [...1],WHI ;wait until high      >Wait high
SETTO:  MOVI  V1,24      ;set 50 step timeout   >Wait low
WLO:    DIBNE [...0],WLO ;wait for falling    >Wait low
TOHI:   DIBEQ [...1],TOLO ;wait for high      >Wait high
        DBNZ  V1,TOHI    ;loop if not timed out >Wait high
        JUMP  WHI        ;timed out, restart   >Restart
TOLO:   DIBEQ [...0],GOTIT ;jump if found events >Wait low
        DBNZ  V1,TOLO    ;loop if not timed out >Wait low
        JUMP  SETTO      ;timed out, restart   >Restart
GOTIT:  ...             ;here for 2 close pulses
```

DISBEQ, DISBNE These instructions test digital input bits 7-0 read by the last **DIBEQ**, **DIBNE** or **WAIT** against a pattern. **DISBEQ** branches on a match. **DISBNE** branches if it does not match.

```
DISBNE [pattern]|Vn|[Vn+off],LB
DISBEQ [pattern]|Vn|[Vn+off],LB
```

pattern This is 8 characters, one for each input bit. The characters can be “0”, “1” and “.” meaning match 0 (TTL low), match 1 (TTL high) or match anything. The bit order in the pattern is [76543210]. You may omit the square brackets, however the **Format** command inserts them.

Vn With a variable the result is: (input BAND Vn(7-0)) BXOR Vn(15-8). A result of 0 is a match, not zero is not a match.

LB The destination of the branch if the input was a match (**DISBEQ**) or not a match (**DISBNE**). This label must exist in the sequence.

This example shows a typical use of this instruction. We want to run a different part of the sequence for three trial types signalled by external equipment that writes the trial type to digital input bits 1 and 0; 00 means no trial, 01, 10 and 11 select trial types 1, 2 and 3.

```
TRWAIT:'W DIBEQ   [.....00],TRWAIT ;Wait for trial >Wait...
        DISBEQ   [.....01],TRIAL1
        DISBEQ   [.....10],TRIAL2
        DISBEQ   [.....11],TRIAL3
```

WAIT The **WAIT** instruction causes the sequence to wait until bits 7-0 of the 1401 digital input match a pattern. The digital input port is sampled once every sequencer clock tick until the pattern is found, or until the sequence is sent elsewhere by a keyboard command. **WAIT** copies digital input bits 7-0 to V56 (**VDigIn**) for use by **DISBEQ** and **DISBNE**. It is usually a good idea to have a display message explaining what you are waiting for.

If you use the template-matched signal in the 1401*plus*, input bits 7-0 may be set as outputs and cannot be used as inputs.

```
WAIT      [pattern]|Vn|[Vn+off]
```

pattern This is the input condition to match before the sequence can continue. It is 8 characters long, one for each input bit. The characters can be “1”, “0” or “.” indicating that the input bit in that position must be a one, a zero or don't care. The bits are given in the order [76543210]. You may omit the square brackets round the pattern if you wish; the **Format** command will insert them.

```
WAIT      [.....1] ;wait for bit 0 set>Wait for bit 0
WAIT      [0.....0] ;wait for 7 and 0 clear>Wait 7&0 low
```

Vn With a variable the result is: (input BAND Vn(7-0)) BXOR Vn(15-8). A result of 0 is required to continue.

This instruction is shorthand for:

```
HERE:     DIBNE   [pattern]|Vn|[Vn+off],HERE
```

DIGIN, BZERO, BNZERO These instructions are obsolete and are provided for backwards compatibility. **DIBNE** and **DIBEQ** are more efficient. **DIGIN** reads digital input bits 7-0 (see page 5-4 for the digital input connections), compares them with a pattern and saves the result. The result can be tested with the two branching instructions **BZERO** and **BNZERO**. The result of the comparisons is preserved until the next **DIGIN**. If you use the template-matched signal in the 1401*plus*, input bits 7-0 may be set as outputs and cannot be used as inputs.

```
DIGIN    [pattern]|Vn|[Vn+off] ;Test input state
BZERO   LB                ;branch to LB if result is zero
BNZERO  LB                ;branch to LB if result non-zero
```

pattern This is 8 characters, one for each input bit in the order [76543210]. The characters can be “.”, “0”, “1” and “c”. For “0” or “1”, the result for that bit is 0 if the input bit was the same, or 1 if it was different. “c” means copy the input bit to the result (this is the same as “0”). The result for “.” is always

zero. You may omit the square brackets round the pattern, however the `Format` command will insert them.

`Vn` With a variable the result is: `(input BAND Vn(7-0)) BXOR Vn(15-8)`. The variable equivalent of `[0.c0110.]` is `190 + 256*12`. The `VarValue` script in the `Scripts` folder calculates variable values equivalent to patterns.

`LB` The destination of the branch if the last `DIGIN` produced a zero (`BZERO`) or non-zero (`BNZERO`) result. This label must exist in the sequence.

```
Loop:   DIGIN  [0.c0110.]           ;assume input is 01101011
        BZERO  GoOn                ;result is 00100110 so no branch
        BNZERO Loop                ;This will branch
GoOn:   ...
```

DAC outputs

The output sequencer supports up to 8 DAC (Digital to Analogue Converter) outputs. The *1401plus* and *Power1401* have four DACs and the *Micro1401* has two. However, the *Power1401* can be expanded with up to 8 DAC outputs. The last value written to DACs 0-7 is stored in variables `V57-V64` (you can also refer to these variables as `VDAC0-VDAC7`). The values are stored as 32-bit numbers with the full 32-bit range corresponding to the full range of the DAC. This high resolution allows us to ramp the DACs smoothly.

Values written to the DACs are expressed in units of your choice. The `SET` or `SDAC` directive determines the conversion between the numbers you supply and the DAC outputs. The standard settings for a system with ± 5 Volts DACs is to set the DAC outputs in Volts.

If you write to a DAC that does not exist, the variable associated with the DAC is set as if the DAC were present. Output to *1401plus* DACs 4-7 is mapped back to DACs 0-3. For the *Micro1401* and *Power1401*, output to DACs that do not exist has no effect.

The *Power1401* DAC 2 and 3 outputs are on the rear panel 37-way Cannon D type Analogue Expansion connector. DAC 2 is on pin 36 and DAC 3 is pin 37. Suitable grounds are on the adjacent pins 18 and 19. If you have a *Power1401* top-box with additional DACs on the front panel, the two DAC output on the rear are mapped to be the two highest numbered DACs. For example, if you have a *2709 Spike2 Top Box*, DACs 2 and 3 are available as BNC connections on the front panel, and the rear panel DACs become DAC 4 on pin 36 and DAC 5 on pin 37.

DAC, ADDAC

The `DAC` instruction can write a value to any of the 8 possible DAC outputs. The `ADDAC` instruction adds a value to the DAC output. The output value changes immediately unless the DAC is in use by the arbitrary waveform output, in which case the result is undefined.

```
DAC      n,expr|Vn|[Vn+off],OptLB
ADDAC    n,expr|Vn|[Vn+off],OptLb
```

`n` The DAC number, in the range 0-7. Variable `57+n` is set to the new DAC value such that the full DAC range spans the full range of the 32-bit variable.

`expr` The value to write to the DAC or the change in the DAC value. The units of this value depend on the `SET` or `SDAC` directive; the standard units are Volts. It is an error to give a value that exceeds the DAC output range.

`Vn` When a variable is used, the full range of the 32-bit variable corresponds to the full range of the DAC. You can use the `VDAC32()` function to load a variable using user-defined DAC units.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example sets DAC 2 to 0 Volts, then ramps it to 4.99 Volts in 1 second using steps of 0.01 Volts. The example also shows how to do the same ramp using variables. You can also use the RAMP instruction to ramp a DAC.

```

      SET    1,1,0           ;1 ms per step, DAC scaled to Volts
'R DAC    2,0               ;Ramp 0 to 5      >Ramping
      MOVI   V1,499         ;499 steps      >Ramping
RAMP1:  ADDAC 2,0.01        ;0.01V increment >Ramping
      DBNZ   V1,RAMP1       ;count increments >Ramping
      HALT                   ;task finished   >Done

      'V MOVI  V3,VDAC32(0)  ;Use variables   >Ramping
      MOVI   V2,VDAC32(0.01) ;increment in V2 >Ramping
      MOVI   V1,499         ;499 steps      >Ramping
      DAC    2,V3           ;set initial value>Ramping
RAMP2:  ADDAC 2,V2          ;add increment   >Ramping
      DBNZ   V1,RAMP2       ;count increments >Ramping
      HALT                   ;task finished   >Done

```

It is a property of the signed integer numbers we use that if you add 1 to the maximum possible positive number, the result is the minimum possible negative number. If you use ADDAC repeatedly to add the same value, eventually you will run off the end of the DAC range and come back in at the other end.

DAC units run from -32768 to +32767. In a ± 5 Volt system with 16-bit DACs, this is -5.0000 to +4.99985 Volts. The DAC unit value for +5 Volts is +32768, but this number does not exist in 16-bit signed integers and wraps around to -32768. Because it often happens that users want to set the DAC to full scale, for the DAC command used with `expr` (not with `Vn`), we change requests to set +32768 units to set 32767 units.

Unlike the digital outputs, the DAC output changes when the instruction runs, not at the next sequencer clock tick. This means that the changes may have a time jitter of a few μ s.

DACn, ADDACn These commands provide backwards compatibility with older Spike2 versions. The DACn and ADDACn commands (with $n = 0, 1, 2$ or 3) set and change the 1401 DAC outputs. `expr` is the new DAC output level, or change in level. Variable `V57+n` is set to the new DAC value such that the full DAC range spans the full range of the 32-bit variable. These commands do not support table use.

```

DACn      expr|Vn,OptLB
ADDACn    expr|Vn,OptLB

```

`expr` The value to assign to DAC n (DACn instruction) or to add to the output level (ADDACn instruction). The units of the DAC values are usually Volts, but can be changed by setting a scale factor with SET or SDAC.

`Vn` With a variable, the values -32768 to 32767 correspond to the full DAC range. You can use the `VDAC16()` function to load a variable using user-defined DAC units. You can also use the `VarValue` script in the `Scripts` folder to calculate variable values. The value saved in `V57+n` is $Vn * 65536$.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

The important difference between these commands and DAC and ADDAC is where a variable is used. The bottom 16-bits of the variable are written to the DAC. In the case of the DAC and ADDAC commands, the upper 16 bits of the variable are written to the DAC.

RAMP This command starts a DAC ramping, with automatic updates on every sequencer step. If the DAC was generating cosine output, the cosine output stops. The DAC ramps from the current value until it reaches a target value, when the DAC cycle flag sets. You can use `WAITC` to test for the end of the ramp. The `RATE` instruction stops a ramp before it reaches the target value. The *1401plus* does not support this instruction and treats it as a NOP.

```
RAMP      n, target|Vn, slope|Vs| [Vs+off]
```

- `n` The DAC number in the range 0-7 for a Power1401 or 0-1 for the Micro1401.
- `target` This is the DAC value at which to end the ramp. The units of the DAC values are usually Volts, but can be changed by setting a scale factor with `SET` or `SDAC`.
- `Vn` When a variable is used for the target, the full range of the 32-bit variable corresponds to the full range of the DAC. You can use the `VDAC32()` function to load a variable using user-defined DAC units.
- `slope` This expression sets the DAC increment per sequencer step. The sign of the value you set here is ignored as the sequencer works out if it must ramp upwards or downwards to achieve the desired target value. If your DAC is calibrated in Volts, to achieve a slope of 1 Volt per second, use `1.0/s (1.0)` for the slope.
- `Vs` You can also set the slope from a variable or by reading it from the table. In this case, the full range of the 32-bit value represents the full range of the DAC. The absolute value of the 32-bit value is used to change the DAC on each step. To get a slope of 1 user unit per second, use `VDAC32(1.0)/s(1.0)` as the value.

This example ramps the DAC 1 from its current level to 1 Volt over 3 seconds, waits 1 second, then ramps it down to 0 Volts over 5 seconds.

```
RAMP      1,1.0,1.0/S(3) ;start with zero size
...      ;other instructions during ramp
WT1:     WAITC  1,WT1      ;wait for ramp to end >Ramp to 1
        DELAY  S(1)-1     ;wait for a second >Wait 1 sec
        RAMP   1,0,1.0/S(5) ;ramp down
WT2:     WAITC  1,WT2      ;wait for ramp >Ramp to 0
```

The `OFFSET` command has an example that uses `RAMP` with cosine waves.

Cosine output control instructions

The sequencer can output cosine waveforms of variable amplitude and frequency through Power1401 DACs 0-7, Micro1401 DACs 0-1, and through *1401plus* DACs 2 and 3. It is an error to output to a DAC that does not exist in your 1401. Attempts to use DACs 0-1 in a *1401plus* are treated as NOPs. When enabled, the cosine value is computed and output every step. The time penalty per step per DAC is around 10 μ s for the *1401plus*, 4 μ s for the micro1401, about 1 μ s for the Power1401 and Micro1401 and less than 1 μ s for the Power1401 mk II. The output is:

output in Volts = $5 A \cos(\theta + \phi) + \text{offset}$

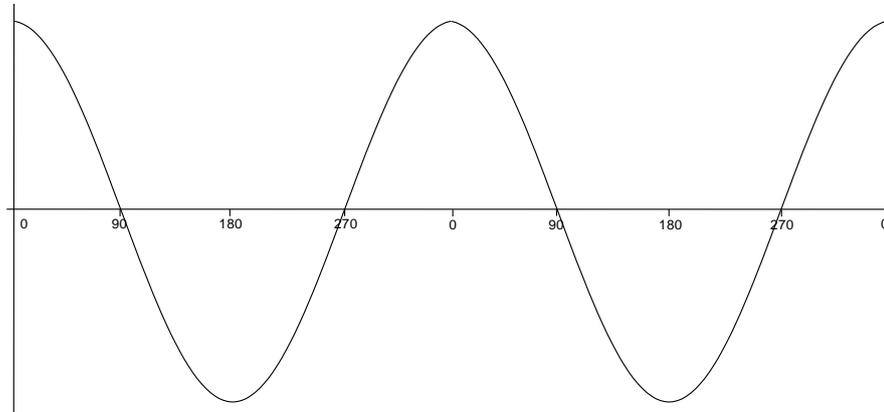
- where `A` is an amplitude scaling factor in the range 0 to 1
- `θ` an angle in the range 0° to 360° that changes each step (set by `ANGLE`)
- `ϕ` is a fixed phase in the range -360° to 360° (set by `PHASE`)
- `offset` A voltage offset set by the `OFFSET` instruction

θ changes every step by $d\theta$. A cycle of the cosine takes $360/d\theta$ steps. You can change the angle increment immediately, or you can delay the change until the next time θ passes through 0° . You can set $d\theta$ in the range 0° up to 360° to an accuracy of about 0.0000001° . With the sequencer running at 1 kHz, you can output frequencies up to 500 Hz with a frequency resolution of around 0.00012 Hz. Ideally the output would be passed through a low pass filter with a corner frequency at one half of the sequencer step rate to smooth out the steps in the cosine wave.

By adjusting ϕ you control the output cosine phase where θ passes through zero. Unless you set the value (PHASE), it is zero and the zero crossing occurs at the peak of the sinusoid. To have the output rising through 0, set the phase to -90 .

Each time θ passes through zero a *new cycle* flag sets. The RAMP, RATEW, RINCW, WAITC and CLRC instructions clear the flag.

Output as a function of $\theta+\phi$



Obsolete commands Before Spike2 version 5.06, the cosine output instructions supported 2 DACs through the Cxxxx and Dxxxx instructions. The Cxxxx family used DAC 1 on the Power1401 and Micro1401 and DAC 3 on the 1401plus. The Dxxxx family used DAC 0 on the Power1401 and Micro1401 and DAC 2 on the 1401plus. The descriptions below include the old forms of the commands, marked *Obsolete*. The DAC numbers next to obsolete commands show the 1401plus DAC number in brackets. There are no plans to remove the old commands, but new sequences should avoid them.

SZ This instruction sets the waveform amplitude. If a wave is playing, the amplitude changes at the next sequencer step. The amplitude is set to 1.0 when sampling starts.

```
SZ      n,expr|Vn|[Vn+off],OptLB      ;DAC n
CSZ     expr|Vn|[Vn+off],OptLB       ;DAC 1(3) - Obsolete
DSZ     expr|Vn|[Vn+off],OptLB       ;DAC 0(2) - Obsolete
```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The cosine amplitude in the range 0 to 1. A cosine with amplitude 1.0 uses the full DAC range.

Vn Variable values 0 to 32768 correspond to amplitudes of 0.0 to 1.0; values outside the range 0 to 32768 cause undefined results.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

SZINC These instructions change the waveform amplitude. The change is added to the current amplitude. If the result exceeds 1.0, it is set to 1.0. If it is less than 0, the result is 0.

```
SZINC  n,expr|Vn|[Vn+off],OptLB      ;DAC n
CSZINC expr|Vn|[Vn+off],OptLB       ;DAC 1(3) - Obsolete
DSZINC expr|Vn|[Vn+off],OptLB       ;DAC 0(2) - Obsolete
```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The change in the waveform scale in the range -1 to 1.

Vn A variable value of 32768 is a scale change of 1.0, -16384 is -0.5 and so on.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

You can gradually increase or decrease the wave amplitude. For example, the following increases the amplitude from zero to full scale (we assume that the waveform is playing):

```

SZ      0,0.0      ;start with zero size
MOVI   V1,100     ;proceed in 1% increments
loop:  SZINC      0,0.01 ;a 1% increase
      DELAY      ms(100)-2 ;show some of the waveform at this size
      DBNZ       V1,loop  ;loop 100 times
    
```

RATE This sets the angle increment in degrees per step, which sets the cosine frequency. If the nominated DAC was ramping, this cancels the ramp. You can stop the cosine output with a rate of 0. Any non-zero value starts the cosine output.

```

RATE      n,expr|Vn|[Vn+off],OptLB ;DAC n
CRATE     expr|Vn|[Vn+off],OptLB   ;DAC 1(3) - Obsolete
DRATE     expr|Vn|[Vn+off],OptLB   ;DAC 0(2) - Obsolete
    
```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The angle increment per step in the range 0.000 up to 180 degrees. The Hz () function calculates the increment required for a frequency.

Vn For a variable, the value 11930465 is an increment of 1 degree. The VHz () function can be used to set a variable value equivalent to an angle in degrees.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example starts cosine output at 10 Hz, runs for 10 seconds, and then stops it. This is then repeated using a variable to produce the same effect:

```

      SET 1,1,0      ;1 ms per step
      'C RATE      0,HZ(10) ;start output at 10 Hz
      DELAY       S(10)-1 ;delay for 10 seconds >Sine wave
X:    'S RATE      0,0      ;stop output
      HALT
      'V MOVI     V1,VHz(10) ;set V1 equivalent of 10 Hz
      RATE       0,V1      ;start at 10 Hz
      DELAY      S(10)-1,X ;delay then goto exit >Sine wave
    
```

RATEW This instruction performs the same function as RATE, except that the change is postponed until the next time θ passes through 0 degrees. RATEW cannot start output; a sinusoid must already be running to pass phase 0. It can stop output, but does not remove the overhead for using cosine output. This instruction clears the new cycle flag (see WAITC).

```

RATEW     n,expr|Vn|[Vn+off]      ;DAC n
CRATEW    expr|Vn|[Vn+off]        ;DAC 1(3) - Obsolete
DRATEW    expr|Vn|[Vn+off]        ;DAC 0(2) - Obsolete
    
```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The angle increment in the range 0.000 to 180 degrees. The Hz () built-in function calculates the increment required for a frequency.

Vn For a variable, the value 11930465 is an increment of 1 degree. The VHz () function can be used to set a variable value equivalent to an angle in degrees.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example starts cosine output at 10 Hz, runs for 1 cycle, changes to 11 Hz for one cycle, then stops:

```

SET 1,1,0           ;1 ms per step
ANGLE 0,0           ;make sure we are at phase 0
RATE 0,Hz(10)      ;start output at 10 Hz
RATEW 0,Hz(11)     ;request 11 Hz next time around
CYCLE10: WAITC 0,CYCLE10 ;wait for the cycle>10Hz
CYCLE11: WAITC 0,CYCLE11 ;wait for the cycle>11Hz
RATE 0,0           ;stop output

```

ANGLE This changes the cosine angular position. It takes effect on the next instruction when the angle increment is added to the value set by this instruction and the result is output.

```

ANGLE n,expr|Vn|[Vn+off],OptLB ;DAC n
CANGLE expr|Vn|[Vn+off],OptLB ;DAC 1(3) - Obsolete
DANGLE expr|Vn|[Vn+off],OptLB ;DAC 0(2) - Obsolete

```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The phase angle to set in the range -360 up to +360.

Vn For a variable, the value 11930465 is a phase of 1 degree (to be precise, 4294967296/360 is a phase of 1 degree). You can use the `VAngle()` function to convert degrees into a suitable value for a variable.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example sets the phase angle to -90 degrees directly, and by using a variable. There is no need to use the `VAngle()` function; we could have set `V1` to -1073741824. However, `VAngle(-90)` is much easier to understand.

```

ANGLE 1,-90        ;set the DAC 1 cosine angle directly
MOVI V1,VAngle(-90)
ANGLE 1,V1         ;set using a variable

```

PHASE This changes the relative phase of the cosine output for the next cosine output. A common use is to change the output from a cosine (maximum value at phase zero) to sine (rising through zero at phase zero).

```

PHASE n,expr|Vn|[Vn+off],OptLB ;DAC n
CPHASE expr|Vn|[Vn+off],OptLB ;DAC 1(3) - Obsolete
DPHASE expr|Vn|[Vn+off],OptLB ;DAC 0(2) - Obsolete

```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The relative phase angle to set in the range -360 up to +360. The relative phase is set to 0 when sampling starts. Set -90 for sinusoidal output.

Vn For a variable, the value 11930465 is a phase of 1 degree (to be precise, 4294967296/360 is a phase of 1 degree). You can use the `VAngle()` function to convert degrees into a suitable value for a variable.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example plays a 1 Hz sinusoidal output (assuming that the output is not running).

```

PHASE 2,-90        ;set the DAC 2 phase angle directly
ANGLE 2,0          ;prepare to start as a sine wave
RATE 2,Hz(1)      ;start the sinusoid

```

OFFSET This changes the cosine output voltage offset for the next cosine output.

```

OFFSET  n,expr|Vn|[Vn+off],OptLB      ;DAC n
COFF    expr|Vn|[Vn+off],OptLB      ;DAC 1(3) - Obsolete
DOFF    expr|Vn|[Vn+off],OptLB      ;DAC 0(2) - Obsolete

```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The offset value for sinusoidal output. The units of this value depend on the SET or SDAC directives; the standard units are Volts. It is an error to give a value that exceeds the DAC output range.

Vn When a variable is used, the full range of the 32-bit variable corresponds to the full range of the DAC. You can use the VDACC32() function to load a variable using user-defined DAC units.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example ramps DAC 0 from 0 to 1 Volt, the runs 5 cycles of a sine wave at 1 Hz, and finally ramps the data back to 0 Volts. This example does not work with a 1401*plus*.

```

          SET      1 1 0          ;1 millisecond per step
          DAC      0,0           ;use DAC 0 for all output
          OFFSET   0,1.0        ;set DAC 0 offset
          SZ       0,0.2        ;1 V sinusoid
          PHASE    0,-90        ;Prepare sinusoid
          ANGLE    0,0          ;set start point
          RAMP     0,1.0,1.0/s(1) ;ramp to 1 Volt in 1 sec
RAMPUP:   WAITC   0,RAMPUP      ;wait for ramp      >Ramp up
          RATE    0,HZ(1)       ;start sinusoid
          DELAY   S(4.9)        ;Sinusoid          >Sine
          RATEW   0,0           ;stop at cycle end
END:      WAITC   0,END         ;wait for end      >Wait end
          RATE    0,0           ;stop now
          RAMP    0,0.0,1.0/S(1) ;ramp to 0 Volt in 1 sec
RAMPDN:   WAITC   0,RAMPDN     ;wait              >Ramp down
          HALT

```

WAITC Each time the phase angle of a cosine passes through 0°, a new cycle flag sets. This flag is also set when a ramp terminates. There is a separate flag for each DAC. This flag is cleared by CLRC, RATEW, RINCW and when tested by WAITC.

```

WAITC   n,LB          ;DAC n
CWAIT   LB           ;DAC 1(3) - Obsolete
DWAIT   LB           ;DAC 0(2) - Obsolete

```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

LB A label to branch to if the new cycle flag is not set. If the flag is set, the sequencer clears the flag and does not branch.

This instruction can produce a pulse at the start (or at least a known time after) the start of each waveform cycle. The following sequence outputs 4 cycles of waveform at different rates on DAC 1, and changes the digital outputs for each cycle.

```

          SZ       1,1.0        ;make sure full size
          ANGLE    1,0.0        ;make sure we start at phase 0
          RATE     1,1.0        ;1 degree per step to start with
          DIGOUT   [00000001]   ;so outside world knows
          RATEW    1,1.2        ;next cycle faster, clear cycle flag
w1:      WAITC    1,w1         ;wait for cycle >1 degree cycle
          DIGOUT   [00000010]   ;announce another cycle
          RATEW    1,1.4        ;next cycle a bit faster
w2:      WAITC    1,w2         ;wait for cycle >1.2 degree cycle
          DIGOUT   [00000011]   ;yet another one

```

```

w3:      RATEW  1,1.6      ;last cycle a bit faster
        WAITC  1,w3       ;wait for cycle >1.4 degree cycle
        DIGOUT [00000100] ;last cycle number
w4:      WAITC  1,w4       ;wait for end >1.6 degree cycle
        RATE   1,0.0      ;stop waveform

```

RINC, RINCW These instructions behave like RATE and RATEW except that they *change* the output rate (angle increment per step) by their argument rather than set it. RINCW clears the new cycle flag.

```

RINC     n,expr|Vn|[Vn+off],OptLB ;DAC n
RINCW   n,expr|Vn|[Vn+off],OptLB ;DAC n
CRINC   expr|Vn|[Vn+off],OptLB   ;DAC 1(3) - Obsolete
CRINCW  expr|Vn|[Vn+off],OptLB   ;DAC 1(3) - Obsolete
DRINC   expr|Vn|[Vn+off],OptLB   ;DAC 0(2) - Obsolete
DRINCW  expr|Vn|[Vn+off],OptLB   ;DAC 0(2) - Obsolete

```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

expr The change in the angle increment per step. You can use the built-in Hz() function to express the change as a frequency.

Vn For a variable, the value 11930465 is a change of 1 degree. You can use the VarValue script in the Scripts folder to calculate variable values.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example starts cosine output at 10 Hz and lets you adjust it from the keyboard.

```

SET      1,1,0      ;1 ms per step
RATE     1,HZ(10)  ;start output at 10 Hz
wt:      JUMP      wt      ;HALT stops all output>P=+1Hz, M=-1Hz
'P RINC  1,HZ(1),wt;1 Hz faster
'M RINC  1,HZ(-1),wt;1 Hz slower

```

These instructions can be used to produce waveforms that change gradually in frequency. The following code generates a linear speed increase every two steps on DAC 1:

```

SZ       1,1.0     ;make sure full size
ANGLE    1,0.0     ;make sure we start at phase 0
RATE     1,1.0     ;1 degree per step to start with
LDCNT1   900      ;in 900 steps of...
loop:    CRINC    0.01 ;...1/100 degrees to...
        DBNZ1    loop ;...10 degrees per step

```

The next example produces 90 cycles, each increasing by 0.1 degrees per step per cycle.

```

SZ       1,1.0     ;make sure full size
ANGLE    1,0.0     ;make sure we start at phase 0
RATE     1,1.0     ;1 degree per step to start with
LDCNT1   90        ;in 90 steps of...
loop:    RINCW    1,0.1 ;...1/10 degrees to...
wait:    WAITC    1,wait ;...(wait for next cycle)...
        DBNZ1    loop ;...10 degrees per step

```

CLRC This instruction clears the cosine output new cycle flag. If you have been running for several cycles and you want to stop the next time phase 0 is crossed use this instruction immediately before using WAITC.

```

CLRC     n,OptLB   ;DAC n
CWCLR    OptLB     ;DAC 1(3) - Obsolete
DWCLR    OptLB     ;DAC 0(2) - Obsolete

```

n The DAC number in the range 0-7 (available DACs depend on the 1401 type).

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example starts a sinusoid and then stops at the next phase 0 crossing after the user requests a stop. Because the sinusoid passes phase 0 in the WAITC instruction and does another step in the RATE 1,0 instruction, we offset the phase by 2 steps. However, this would cause the start of the sinusoid to be 2 steps wrong, so we change the start angle to match.

```
'G PHASE 1,-2*Hz(2) ; compenstate for ending
   ANGLE 1,2*Hz(2) ; so we start in correct place
   RATE 1,Hz(2) ; 2 Hz output
HERE: JUMP HERE ; output is running >Running
   'S CLRC 1 ; Stop output
WT: WAITC 1,WT ; wait for cycle end>Waiting
   RATE 1,0,HERE ; stop and then idle
```

General control These instructions do not change any outputs or read data from any inputs. They provide the framework of loops, branches and delays used by the other instructions.

DELAY The DELAY instruction occupies one clock tick plus the number of extra ticks set by the argument. It produces simple delays of 1 to more than 4,000,000,000 sequencer steps.

```
DELAY expr|Vn|[Vn+off],OptLB
```

expr The extra sequencer clock ticks to delay in the range 0 to 4294967295. The s(), ms() and us() built-in functions convert a delay in seconds, milliseconds or microseconds into sequencer steps.

Vn Variable or table index from which to read the number of extra clock ticks.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

This example uses display messages to tell the user what the sequence is doing.

```
SET 1.00,1,0 ;run with 1 millisecond clock ticks
DELAY 2999 ;wait 2999+1 milliseconds>3 second delay
DELAY s(3)-1 ;3 seconds -1 tick delay >3 second delay
DELAY V1,LB ;wait V1+1 ms, branch >variable delay
DELAY [V1+9] ;V1+9 is table index >table delay
```

DBNZ DBNZ (Decrement and Branch if Not Zero) subtracts 1 from a variable and branches to a label unless the counter is zero. It is used for building loops.

```
DBNZ Vn, LB
```

Vn The variable to decrement and test for zero.

LB Instruction to go to next if the result of the decrement is not zero.

DBNZ is often used with MOVI to set up loops, for example:

```
MOVI V2,1000 ;set times to loop
WT: DIGOUT [00000000] ;set all digital outputs low
   DIGOUT [11111111] ;set them all high
   DBNZ V2,WT ;loop 1000 times
```

LDCNTn, DBNZn These obsolete instructions use variables V61 to V64 as counters 1 to 4. New sequences should use the MOVI and DBNZ instructions that can use any variable as a counter. The

LDCNTn instruction loads one of the four counters. DBNZn (Decrement and Branch if Not Zero) decrements a counter and branches to a label while the counter is not zero.

```
LDCNTn  expr|Vn
DBNZn   LB
```

- n The counter number to load or decrement in the range 1 to 4. Counters 1 to 4 are emulated by variables V61 to V64.
- expr The 32-bit integer value to load into the counter.
- Vn When a variable is used the counter is set to the 32-bit variable.
- LB Instruction to go to next if the result of the decrement is not zero

CALL, CALLV, CALLn, RETURN

These instructions run a labelled part of a sequence and return. CALL, CALLV and CALLn (obsolete) save the next step number to a *return* stack and jump to the labelled instruction. The RETURN instruction removes the top step number from the return stack and jumps to it. CALLV also sets a variable to a constant. CALLn (n=1 to 4) is obsolete and sets the value of the variables that emulate the four counters V61 to V64.

```
CALL    LB                ; use LB as a subroutine
CALLV   LB,Vn,expr       ; Vn = expr, then call LB
CALLn   LB,expr          ; V60+n = expr (n=1, 2, 3 or 4)
RETURN  LB               ; return to step after last CALL
```

- LB The next instruction to run. The CALLED section should end with a RETURN.
- Vn CALLV copies the value of expr to this variable. CALL1 sets V61, CALL2 sets V62, CALL3 sets V63 and CALL4 sets V64.
- expr A 32-bit integer constant that is copied to a variable. For CALLn only, if expr is 0, the variable is set to 256 to be compatible with previous versions.

You can use CALL inside a CALLED subroutine. This is known as a *nested* CALL. If you call a subroutine from inside itself, this is known as a *recursive* CALL. The return stack has room for 64 return addresses. If you use more than this, the oldest return address is overwritten, so your sequence will not behave as you expect.

This example generates different pulse widths from DAC 0. The sequence is written to be independent of the sequencer rate (it must be high enough so that the widths are possible). In this case the rate is set to 5 kHz. The example sets DAC 0 to zero, then pulses for 20 milliseconds twice, once using CALL and once using CALLV. Then after a delay, there is a 50 millisecond pulse.

```
SET     0.2,1,0          ; Run at 5 kHz, normal DAC scale
DAC     0,0              ; make sure DAC0 is zero
MOVI    V3,ms(20)-2     ; these two instructions...
CALL    PUL             ; ...have the same effect as...
CALLV   PUL,V3,ms(20)-2; ...this one. 20 ms pulse
DELAY   s(1)-1         ; wait 1 second, then...
CALLV   PUL,V3,ms(50)-2; ...a 50 ms pulse
HALT    PUL             ; So we don't fall into PUL routine
PUL:    DAC     0,1      ; set DAC value
        DELAY   V3      ; wait for time set
        DAC     0,0      ; set DAC back to zero
        RETURN  PUL     ; back to the caller
```

CALL/CALLV and RETURN let you reuse a block of instructions. This can make sequences much easier to understand and maintain. The disadvantage is the additional steps for the CALL and RETURN. If you need to set a variable, use CALLV and there is only the overhead of the RETURN instruction.

JUMP The `JUMP` instruction transfers control unconditionally to the instruction at the label. Many instructions allow the use of an optional label to set the next instruction, so you can often avoid the need for this instruction. You can also jump using the contents of a register as the destination, or relative to a label (LB):

```
JUMP  LB                ; Jump to label
JUMP  (Vn),OptLB       ; use MOVI Vn,LB to set target
JUMP  LB(Vn),OptLB     ; Jump to instruction LB+Vn
```

(Vn) The value of variable `Vn` sets the instruction number to jump to.

LB(Vn) Jump to the instruction given by label `LB` plus the contents of `Vn`.

OptLB An optional label to jump to if (Vn) or LB(Vn) is not an instruction number.

HALT The `HALT` instruction stops the output sequence and removes all overhead associated with it. It does not stop the sequencer clock, which continues to run. Any cosine output will stop, but will restart when the sequence restarts. To restart the sequencer, press a key associated with a sequence step or click a key in the sequencer control panel. If you associate a display string with this instruction, it appears in the sequencer control panel.

```
HALT                                     >Press X when ready
```

NOP The `NOP` instruction (No Operation) does nothing except use up one sequencer clock tick. It can be thought of as the equivalent of `DELAY 0`.

Variable arithmetic These instructions perform basic mathematical functions while a sequence runs. You can also compare variables and branch on the result.

Compare variable These instructions compare a variable with a variable or a 32-bit expression or a table entry and branch on the result. All comparisons are of signed 32-bit integers.

```
Bxx   Vn,Vm,LB         ;compare with a variable
Bxx   Vn,expr,LB      ;compare with a constant
Bxx   Vn,[Vm+off],LB  ;compare with a table entry
```

xx This is the branch condition. The `xx` stands for: `GT=Greater Than`, `GE=Greater or Equal`, `EQ=Equal`, `LE=Less than or Equal`, `LT=Less Than`, `NE=Not Equal`.

Vn The variable to compare with the next argument.

Vm A variable to compare `Vn` with or table index variable.

expr A 32-bit integer constant to compare `Vn` with.

This example collects the latest data value from channel 1 (assumed to be a waveform), waits for it to be in a preset range for 1 second, then outputs a pulse to a digital output bit.

```
START: CHAN  V1,1          ; get channel 1 data
       BGT   V1,4000,START ; if above upper limit, wait
       BLT   V1,0,START   ; if too low, wait
IN:    MOVI  V2,S(1)/4     ; timeout, 4 instructions/loop
INLOOP: CHAN  V1,1        ; to check if still inside
       BGT   V1,4000,START ; if above upper limit, wait
       BLT   V1,0,START   ; if too low, wait
       DBNZ  V2,INLOOP    ; see if done yet
REWARD: DIGOUT [...1]    ; Task done OK
       DELAY S(1)        ; leave bit set for 1 second
```

```
DIGOUT [.....0] ; clear done bit
... ; next task...
```

We want the data to be in range for one second. There are 4 instructions in the loop that tests this, so we set to the loop to run for the number of steps in a second divided by 4. For this to work correctly, the sequencer must be running fast enough so that 4 steps are no longer than the sample interval for the waveform channel.

MOVI This instruction moves an integer constant into a variable. `MOVIL` is an obsolete instruction that does exactly the same thing. The syntax is:

```
MOVI Vn,expr,OptLB ; Vn = expr
```

`Vn` A variable to hold the value of `expr`.

`expr` An expression that is evaluated as a 32-bit integer.

`OptLB` If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

`MOVI` is not the same as the `VAR` directive. The `VAR` directive sets the value of a variable when the sequence is copied to the 1401 and does not occupy a step. The `MOVI` instruction is part of the sequence and set the value of the variable each time the instruction is used.

MOV, NEG, ABS The `MOV` instruction sets a variable to the value of another with the option of adding a 32-bit number and dividing by a power of two). `ABS` takes the absolute value first. `NEG` is identical to `MOV` except that the source variable is negated first. The syntax is:

```
MOV Va,Vb,expr,shift ; Va = (Vb + expr) >> shift
NEG Va,Vb,expr,shift ; Va = (-Vb + expr) >> shift
ABS Va,Vb,expr,shift ; Va = (|Vb| + expr) >> shift
```

`Va` A variable to hold the result. It can be the same as `Vb`.

`Vb` A variable used to calculate the result. It can be the same variable as `Va`.

`expr` An optional expression that is evaluated as a 32-bit integer. If this argument is omitted, it is treated as 0.

`shift` An optional argument in the range 0 to 31, set to 0 if omitted, that sets the number of times to divide the result by 2.

The following examples assume that `V3` holds 1000:

```
VAR V6,Result
MOV V1,V3 ; set V1 to 1000
NEG V1,V3 ; set V1 to -1000
MOV V1,V3,-8 ; set V1 to 992
NEG Result,V3,0,4 ; set V6 to -63
MOV Result,V3,4,1 ; set V6 to 502
```

ADDI This instruction adds a 32-bit integer constant to a variable. `ADDIL` is an obsolete instruction that does exactly the same. There is no `SUBI` as you can add a negative number. The syntax is:

```
ADDI Vn,expr,OptLB ; Vn = Vn + expr
```

`Vn` A variable to hold the result of `Vn + expr`.

`expr` An expression that is evaluated as a 32-bit integer.

`OptLB` If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

The following examples assume that v1 holds -1000:

```
VAR      V1,Result=-1000
ADDI     Result,1000      ; set V1 to 0
ADDI     V1,-4000        ; set V1 to -5000
```

ADD, SUB The ADD instruction adds one variable to another. The SUB instruction subtracts one variable from another. In both cases you can optionally add a 32-bit integer constant and optionally divide the result by a power of two. The syntax is:

```
ADD      Va,Vb,expr,shift ; Va = (Va + Vb + expr) >> shift
SUB      Va,Vb,expr,shift ; Va = (Va - Vb + expr) >> shift
```

Va A variable to hold the result. It can be the same as Vb.

Vb A variable to add or subtract.

expr An optional expression evaluated as a 32-bit integer. If omitted, 0 is used.

shift An optional argument in the range 0 to 31, set to 0 if omitted, that sets the number of times to divide the result by 2.

The following examples assume that v1 holds -1000, v3 holds 1000, v6 holds 100:

```
VAR      V6,Result=100
ADD      V1,V3          ; V1 = 0 (-1000 + 1000 + 0)
SUB      V1,V3          ; V1 = -1000 (0 - 1000 + 0)
ADD      V1,V3,-8       ; V1 = -8 (-1000 + 1000 - 8)
SUB      Result,V3,0,2  ; V6 = -225 (100 - 1000 + 0)/4
ADD      Result,V3,4,1  ; V6 = 389 (-225 + 1000 + 4)/2
```

MUL, MULI MUL multiplies a variable by another variable, then optionally adds a 32-bit integer constant and divides the result by a power of two. MULI multiplies a variable by a 32-bit integer constant and divides the result by a power of two.

```
MUL      Va,Vb,expr,shift ; Va = ((Va*Vb)+expr) >> shift
MULI     Va,expr,shift    ; Va = (Va*expr) >> shift
```

Va A variable to hold the result. It can be the same as Vb.

Vb A variable used to calculate the result.

expr An expression that is evaluated as a 32-bit integer. It is optional for MUL and required for MULI. If this argument is omitted, it is treated as 0.

shift An optional argument in the range 0 to 31, set to 0 if omitted, that sets the number of times to divide the result by 2.

The following examples assume that v1 holds -10 and v3 holds 10:

```
MULI     V1,10          ; V1 = -100 (-10 * 10)
MUL      V1,V3,-8       ; V1 = -992 (-10 * 10 -8)
```

DIV, RECIP DIV and RECIP divide variables and were added at version 5.08. They take around 1 μ s in the Power1401, 3 in a Micro1401, 10 in a micro1401 and 5 in a 1401*plus*.

```
DIV      Va,Vb          ; Va = Va / Vb
RECIP    Va,expr        ; Va = expr / Va
```

If the numerator is 0, the result is 0. If the denominator is 0, the result is 2147483647 if the numerator is greater than 0 and -2147483648 if it is negative. The 1401*plus* truncates all results downwards, all other 1401s truncate towards 0. So, 7/3 or -7/-3 is 2 in all 1401s, but -7/3 or 7/-3 is -2 except in a 1401*plus*, where it is -3.

Variable logic These instructions perform bitwise logical functions while a sequence runs.

AND, ANDI These instructions bitwise AND a variable with a variable or a 32-bit expression. A bitwise AND means that each bit of the 32-bit result is 1 if both corresponding source bits are 1, otherwise the result bit is 0. For example, 3 AND 1 is 1, 0x55 AND 0xAA is 0.

```
AND    Va,Vb          ; Va = Va AND Vb
ANDI   Va,Vb,expr    ; Va = Vb AND expr
```

Va The variable to hold the result.

Vb A variable to AND with Va or with the expression.

expr A 32-bit integer constant to AND with Va.

This example waits for the digital input to have bit 4 set, then branches based on the digital input value (placed in VDigIn or V56 by WAIT).

```
WAIT   [...1....]    ; wait for bit 4 set
ANDI   V1,VDigIn,0x0f ; isolate bits 0..3 (value 0-15)
JUMP   ACTION(V1)    ; branch based on the result
ACTION: JUMP  ACT0    ; action for value 0
        JUMP  ACT1    ; action for value 1
        ...
        JUMP  ACT15   ; action for value 15
```

OR, ORI These instructions bitwise OR a variable with a variable or a 32-bit expression. A bitwise OR means that each bit of the 32-bit result is 1 if either corresponding source bit is 1, otherwise the result bit is 0. For example, 3 OR 1 is 3, 0x55 OR 0xAA is 0xff.

```
OR     Va,Vb          ; Va = Va OR Vb
ORI    Va,Vb,expr    ; Va = Vb OR expr
```

Va The variable to hold the result.

Vb A variable to OR with Va or with the expression.

expr A 32-bit integer constant to OR with Va.

XOR, XORI These instructions bitwise exclusive OR a variable with a variable or a 32-bit expression. A bitwise exclusive OR means that each bit of the 32-bit result is 1 if the corresponding source bits differ, otherwise the result bit is 0. For example, 3 XOR 1 is 2, 0x55 XOR 0xAA is 0xff.

```
XOR    Va,Vb          ; Va = Va XOR Vb
XORI   Va,Vb,expr    ; Va = Vb XOR expr
```

Va The variable to hold the result.

Vb A variable to XOR with Va or with the expression.

expr A 32-bit integer constant to XOR with Va.

Table access Tables are declared with the TABSZ directive and can be populated with data using the TABDAT directive. Most access to tables is through the [Vn+off] method, but there are also instructions for loading and storing a variable in a table and for incrementing or decrementing a variable used as a pointer into the table.

TABLD, TABST These two instructions load a variable from the table and store a variable into the table. Many instructions can load arguments from the table, so `TABLD` is not often required.

```
TABLD  Vm, [Vn+off], OptLB ; load Vm from the table
TABST  Vm, [Vn+off], OptLB ; store Vm into the table
```

- `Vm` The variable to load from the table or store into the table.
- `+off` An optional expression that evaluates to an integer in the range -1000000 to 1000000 . If omitted, 0 is used.
- `Vn` The variable value plus the offset is used as a table index. If the index lies in the table, `Vm` is loaded from the table or stored in the table at the index. If the index is not in the table, `TABLD` copies 0 to `Vm` and `TABST` does nothing.
- `OptLB` If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

TABADD, TABSUB These two instructions add a table value to a variable or subtract a table value from a variable. These instructions were added at version 7.00.

```
TABADD  Vm, [Vn+off], OptLB ; add table value to Vm
TABSUB  Vm, [Vn+off], OptLB ; subtract value from Vm
```

- `Vm` The variable to add data to or subtract it from.
- `+off` An optional expression that evaluates to an integer in the range -1000000 to 1000000 . If omitted, 0 is used.
- `Vn` The variable value plus the offset is used as a table index. If the index lies in the table, `Vm` is changed. If the index is not in the table, the instruction does nothing.
- `OptLB` If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

TABINC This instruction adds a constant to a variable and detects if the result is a valid table index. If it is a valid index, the instruction branches. If it is not, the result is reduced by the table size if it is positive and is increased by the table size if it is negative, and the instruction does not branch. This gives you an efficient way to work through the table.

```
TABINC  Vn, expr, OptLB
```

- `Vn` This variable is assumed to hold a valid table index.
- `expr` This expression evaluates to a positive or negative number that is added to `Vn`.
- `OptLB` If present, branch to this label if `Vn+expr` is a valid table index.

For example, the following codes plays pulses through DAC 0 based on data in the table. The table data holds groups of three items, holding the time for the DAC to stay at 0, the DAC amplitude and the time to stay at the amplitude. Some example table data is given, but the data could also be set with the `SampleSeqTable()` script command.

```

SET      0.100 1 0           ;run at 10 kHz
TABSZ    12                  ;4 sets of 3 items
TABDAT   ms(50)-2, VDAC32(1), ms(50)-3
TABDAT   ms(100)-2, VDAC32(1.3), ms(70)-3
TABDAT   ms(200)-2, VDAC32(1.5), ms(90)-3
TABDAT   ms(400)-2, VDAC32(1.9), ms(110)-3
'G MOVI  V1, 0               ;use V1 as the table pointer
LOOP:    DAC      0, 0        ;strt with the DAC low
         DELAY    [V1]        ;wait for first period>Low
         DAC      0, [V1+1]    ;get the DAC value
         DELAY    [V1+2]      ;wait for second period>High
         TABINC   V1, 3, LOOP
         DAC      0, 0        ;tidy up the dac
```

Access to data capture

Most activities in the sequencer are independent of the sampling process. However, there are times when you need to know the value of a channel to decide what to do next. The `CHAN` command gives you the latest waveform value or number of events on a channel. The `TICKS` command tells you the current time in terms of the sampling clock ticks.

The sequencer can also send information in the other direction. If the digital marker channel is active you can force it to record the current digital input state with `REPORT`, or you can force it to record a marker of your own choosing with `MARK`.

CHAN This instruction gives the output sequencer access to sampled data on a waveform channel and to the number of recorded events on all other channel types. You can also use this command to get the most recent value written to the DAC outputs. The variable value is 0 if the channel is not being sampled.

```
CHAN    Vn,chn          ; Vn = ChanData(chn)
```

`chn` The channel number is 1 to 100 for sampled channels or 0 to -7 for the last value on DAC 0 to 7. The result is the most recent data available. For a slow waveform channel this could be a long time in the past. In triggered sampling mode, waveform data is available between triggers.

Waveform and DAC data are treated as 16-bit signed values from -32768 to 32767 for version 3 compatibility. You also have access to DAC values as 32-bit data in variables V57 to V64 (VDAC0 to VDAC7) without the need to use the `CHAN` instruction.

This example waits for a signal to cross 0.05 volts and produces a pulse. We assume that channel 1 is a waveform.

```
SET      0.100 1 0      ;run at 10 kHz
VAR      V1,level=VDAC16(0.05) ;level to cross
VAR      V2,data        ;to hold the last data
VAR      V3,low=0       ;some sort of hysteresis level
DIGOUT   [00000000]    ;set all dig outs low
BELOW:   CHAN  data,1    ;read latest data >wait below
        BGT   data,low,below ;wait for below >wait below
ABOVE:   CHAN  data,1    ;read latest data >wait above
        BLE  data,level,above ;wait for above >wait above
DIGOUT   [.....1]     ;pulse output...
DIGOUT   [.....0],below ;...wait for below
```

TICKS This instruction sets a variable to the current sampling time in Spike2 time units (microseconds per time unit set in the sampling configuration) and adds an expression or 0 if `expr` is omitted. The `sTick()`, `msTick()` and `usTick()` expression functions can be used to make the sequence independent of the microseconds per time unit value.

```
TICKS   Vn,expr        ; Vn = Spike2 time + expr
```

This can be used with the `CHAN` command and variable related branches to check the timing of external pulses. The sequencer runs under interrupt, and competes for time with other interrupt driven processes in the 1401 interface. This causes some “jitter” in the timing. The jitter for a Micro1401 or Power1401 is typically only a few microseconds. For a 1401*plus*, it can be a few tens of microseconds, depending on other 1401 activity.

REPORT, MARK

The `REPORT` instruction records a digital marker (if the digital marker channel is enabled) as if there were an external digital marker trigger pulse. The `MARK` instruction does the same, except it takes the argument as the value to record. `REPORT` has no arguments.

```
REPORT   OptLB
MARK     expr|Vn|[Vn+off],OptLB
```

expr The argument should have a value in the range 0 to 255. If a variable or table is used, the bottom 8 bits of the value are used.

OptLB If this optional label is present it sets the next instruction to run, otherwise the next sequential instruction runs.

```

WAIT      [.....1]      >Waiting for bit 0
REPORT                   ;save a marker when this is set
MARK      12             ;set code 12 as a digital marker
    
```

Randomisation

These functions use a pseudo-random number generator. The generator is seeded by a number that is based on the length of time that the 1401 has been switched on.

BRAND

BRAND branches with a probability set by the argument or by a variable. This could be used when several different stimuli are required, but in a random sequence. BRANDV is an obsolete name for the same instruction.

```
BRAND    LB,expr|Vn|[Vn+off]
```

LB Where to go if the branch is taken

expr This is the probability of branching in the range 0 up to (but not including) 1.

Vn When a variable or table entry is used for a branch, the value is treated as a 32-bit unsigned number; 0 means a probability of zero and 4294967295 (the largest 32-bit unsigned number) means a probability of 0.999999998.

```
BRAND    LB,0.5           ;branch with 50% probability
```

To produce a multiple way random branch you use more BRAND instructions. A three way equal probability branch to LA, LB and LC can be coded:

```

BRAND    LA,0.33333   ;Split the first route with p=1/3
BRAND    LB,0.5       ;0.6667 to here * 0.5 is 0.3334 (1/3)
LC:      ...           ;If neither of the above, comes here
    
```

The following shows the sequence for a five-way branch with equal probabilities:

```

BRAND    LA,0.2       ;5 way, LA probability is 0.2 (1/5)
BRAND    FX,0.5       ;Probability to here=0.8, so to FX=0.4
BRAND    LB,0.5       ;Probability to here=0.4, so to LB=0.2
LC:      ...           ;Probability to here=0.2
FX:      BRAND    LD,0.5 ;Probability to here=0.4, so to LD=0.2
LE:      ...           ;Probability to here=0.2
    
```

The best technique is to reduce the branches to a power of two as soon as possible. Case 1 of the five-way branch is split off (probability of 0.2), leaving 4 ways. The 4 ways are split with a probability of 0.5 (0.4 for each division) then the last two routes are split, again with a probability of 0.5 (0.2 for each division).

Poisson process

In a Poisson process, the probability of something happening per time interval is constant. You can generate a delay with a Poisson statistic by:

```
POISSON: BRAND   POISSON,prob ; poisson delay
```

The probability is given by $prob = 1.0 - 1.0/(mDelay * S(1))$, where *mDelay* is the mean delay required in seconds and *S(1)* is the built in function that tells us how many steps there are per second. If you would rather express this in terms of a rate, then $prob = 1.0 - rate/S(1)$, where *rate* is the expected rate in Hz.

```

TENHZ:   BRAND    TENHZ,1.0-10/S(1) ;10 Hz mean rate
          DIGOUT [.....1]         ;set output high
          DIGOUT [.....0],TENHZ   ;set output low, goto TENHZ
    
```

This example generates a digital output that pulses to produce an approximation to a Poisson distributed pulse train with a mean frequency of 10 Hz. The approximation improves the shorter the step time. The mean interval between pulses is 100 milliseconds plus the time for 2 steps and the shortest gap between pulses is 3 sequencer steps.

Scripts and variables From a script you can set sequencer variables as 32-bit signed integers. For the range 2147483648 to 4294967295 we must use negative numbers. This script example shows you how to convert a probability into a variable value and pass it to the sequencer:

```
Proc SetBrandVar(prob, v%) 'prob is probability, v% is variable
prob *= 4294967296.0;      'range 0-4294967296 is 0 to 1.0
if (prob > 4294967295.0 ) then prob := 4294967295.0 endif;
if prob > 2147483647 then prob -= 4294967296.0 endif;
if prob < -2147483647 then prob := -2147483647 endif;
SampleSeqVar(v%, prob);
end;
```

MOVRND This instruction generates a random number in the range 0 to a power of 2 minus 1, then adds an integer constant to it and stores the result in a variable.

```
MOVRND Vn,bits,expr
```

Vn The variable to hold the result.

bits The number of random bits to generate in the range 1 to 32. The generated random bits fill the variable starting at the least significant bit. Bits above the highest numbered generated bit are set to 0.

expr An optional expression that evaluates to a 32-bit integer number, that is added to the random bits. If this is omitted, nothing is added.

Expressed in terms of the script language, the random number is one of the numbers in the range $expr$ to $expr + \text{Pow}(2, \text{bits}) - 1$. For example, `MOVRND V61, 8, 1` emulates the obsolete `LD1RAN` instruction that generates a random number in the range 1 to 256.

The following code fragment implements a random delay of between 1 and 2.024 seconds (assuming a 1 millisecond clock).

```
MOVRND V1,10,998 ;load V1 0 with (0 to 1023) + 998
DELAY V1 ;this uses 999 to 2023 steps
```

LD1RAN This obsolete instruction loads counter 1 (V61) with a pseudo-random number in the range 1 to 256. It is implemented as `MOVRND V61, 8, 1`.

Arbitrary waveform output

In addition to generating voltage pulses, ramps and cosine waves through the DACs, Spike2 can play arbitrary waveforms. The sequencer can start waveform output, test it and branch on the result, stop output or set one more output cycle. Each waveform has an associated code (often a keyboard character). The sequencer uses the code to identify the arbitrary wave to replay. The waveforms are stored in 1401 memory and can be updated during sampling with the `PlayWaveCopy()` script command. See the *Sampling data* chapter of the *Spike2 for Windows* manual for more about arbitrary waveform output.

WAVEGO The `WAVEGO` instruction starts output from a play wave area, or prepares the output for an external trigger. Starting output can take more time than we want to allocate to a sequencer step so `WAVEGO` sets a flag and output starts as soon as the 1401 has free time (usually within a millisecond). See `WAVEST` for a precisely timed start to output.

```
WAVEGO code{, flags}
```

code This is either a single character standing for itself, or a two digit hexadecimal code. This is the code of the area to be played.

flags These are optional single character flags. The flags are not case sensitive. Use **T** for triggered waveform output. Use **w** to make the sequencer wait at this step until the 1401 has prepared the hardware to play. For example:

```
WAVEGO X           ;area X, no wait, no trigger
WAVEGO 23,T        ;area coded hexadecimal 23, triggered
WAVEGO 0,WT        ;area 0, wait until trigger armed
WAVEGO 1,W         ;area 1, wait until play started
```

If you need to know when the output started, use the **WAVEBR T** option. If you need to know that the request to start the playing operation has been honoured, but you do not want to hang up the sequencer with the **w** option, use the **WAVEBR W** option.

WAVEGO cancels existing output just before the new area starts to play. If you use **WAVEBR** after the play request, unless you use the **WAVEGO** or **WAVEBR W** option to be certain that the new area is active, the **WAVEBR** result may be based on the previous area.

WAVEBR The **WAVEBR** instruction tests the state of the waveform output and branches on the result. No branch occurs if there is no output running or requested.

```
WAVEBR LB, flag
```

LB Label to branch to if the condition set by the flag is not met.

flag An optional single character flag to specify the branch condition:
w branch until a **WAVEGO** request without the **w** flag is complete.
C branch until the play wave area or cycle count changes.
A branch until the play wave area changes.
S branch until the current output stops.
T branch until output started with **WAVEGO** begins to play.

The following sequence tracks the output when we have two play areas labelled 0 and 1. Area 0 is set to play 10 times and is linked to area 1. The sequence below will track the changes. Play wave DAC output happens one play wave clock tick after the output is changed, so the sequencer can know that a DAC output is about to change.

```
WAVEGO 0,WT        ; area 0, wait for armed
LDCNT1 5           ; load counter 1 with 5
WT:  WAVEBR WT,T   ; wait for external trigger>Trigger?
W5:  WAVEBR W5,C   ; wait for cycle >Waiting for cycle
      DBNZ1 W5     ; do this 5 times >Waiting for cycle
WA:  WAVEBR WA,A   ; wait for area >Wait for area
WE:  WAVEBR WE,S   ; wait for end >Wait for end
```

The **WAVEGO** command requests a triggered start and waits until the trigger is armed before moving on. It then waits for an external trigger at the **WT** label. Next the sequence tracks the end of 5 output cycles. At label **WA** the sequence waits for the area to change and finally the sequence waits for output to stop.

WAVEST The **WAVEST** instruction can start output that is waiting for a trigger and stop output that is playing, either instantly, or after the current cycle ends (see the *Sampling data* chapter for more details of triggering, including connections).

```
WAVEST flag
```

flag This are optional single character flag and specifies the action to take:
T Trigger a waveform that is waiting for a triggered start.

- S Stop output immediately, no link to the next area.
- C Play to the end of the current cycle, then end this area. If there is another area linked it will play.

The following code starts output with an internal trigger and then stops it

```
WAVEGO X,TW          ; arm area X for trigger
WAVEST T             ; trigger the area
DELAY 1000           ; wait
WAVEST S             ; stop output now
```

Compatibility with previous versions

Sequences from all previous versions still work, with the following provisos on sequences written for Spike2 version 3 and earlier:

1. Variables and labels may not share names nor be the same as instruction codes.
2. DELAY 0 no longer hangs up for a long time. It now has the same effect as NOP.
3. If a version 3 sequence relied on looping back to the start after 256 steps, this no longer works as all sequences have a HALT instruction automatically added to the end.
4. In version 3, BRAND set the probability to an accuracy of one part in 256. It now sets the probability to a very high accuracy. This might affect your results.
5. DACn and ADDACn now insist that any value you use lies within the range of the DAC.
6. ADDACn previously expressed your increment as a 16-bit integer. It now generates a much more accurate 32-bit integer. However, this means that ramps generated with ADDACn may have a different slope (probably much closer to the intended slope).

If you want to write a sequence that will run on previous versions of Spike2, either refer to the manual for the previous version, or write the sequence using the previous version to be certain of complete compatibility. Use of any of the following features will prevent the sequence working on version 6 or earlier:

1. Use of more than 1023 sequencer instructions or more than 64 variables.
2. Use of the ASz () or VSz () expressions.
3. Use of a label in an expression
4. Use of the JUMP (Vn) or JUMP LB(Vn) instructions
5. Use of new instructions: TABADD, TABSUB, ABS, AND, ANDI, OR, ORI, XOR, XORI
6. Use of #include to include other files.
7. Use of the = directive to define a constant.

Sequencer compiler error messages

When you use the Current, Format or Compile buttons in the sequence editor, Spike2 checks the sequence and displays the result of the compilation or format operation in the message bar at the top of the window. If there is an error, the offending line is marked by changing the background colour. The messages report either successful operation or the cause of the problem. When compiling with any #include files, if a problem occurs in an included file, that file opens to highlight the problem.

File menu

The File menu is used for operations that are mainly associated with documents (opening, closing, importing and creating), configuration files and with printing. The final command in the File menu is also your route out of Spike2.

New File



This command creates a new Spike2 data file for sampling, an output sequence file, a script file, text file or an XY file. The command opens a new file dialog box in which you choose the type of file to create. You can activate this command from the menu and from the toolbar.



You can make files of five types: data, output sequencer (pulse) files, script files, text files and XY files. Double-click a file type or select one and click OK; Spike2 will open a window of the specified type. Each file type has its own file name extension.

Data Document A sampling window opens plus additional windows set by the sampling configuration (see the *Sampling data* chapter for details). Data documents are not stored in memory, unlike most new files, but are kept on disk. Until they are saved after sampling they are temporary files in the directory set in the Edit menu Preferences. The file name extension is `.smr`. Only one sampling window can be open at a time.

Result Document These hold result views (file extension `.srf`). Result views are created as the result of an analysis, or from the script language, not with the File menu New command.

Sequencer Document A new window opens in which you can type, edit and compile an output sequence (see the *Data output during sampling* chapter for details). The file name extension is `.pls`.

Text Document Text files can be used to take notes, build reports and to cut and paste text between other windows and applications. The file name extension is `.txt`.

Script Document A script editor window opens in which a new script can be written, run and debugged. The file name extension is `.s2s`.

XY Document XY windows are used to draw user-defined graphs with a wide variety of line and point styles. The Windows file name extension is `.sxy`. They can be generated by scripts, or from the Analysis menu for trend plots.

Other types used by Spike2 Spike2 also uses other file types. Files with an extension ending in `x` are XML format; we also read (but no longer write) older versions of these without the `x`.

Resource files Spike2 creates files with the extension `.s2rx`. These are associated with data files of the same name, but with the extension `.smr`. They hold configuration information so that Spike2 can restore the display. These files are not essential to Spike2 and if deleted, the associated data file is not damaged in any way.

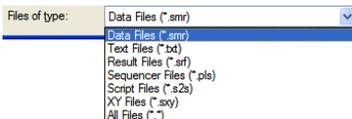
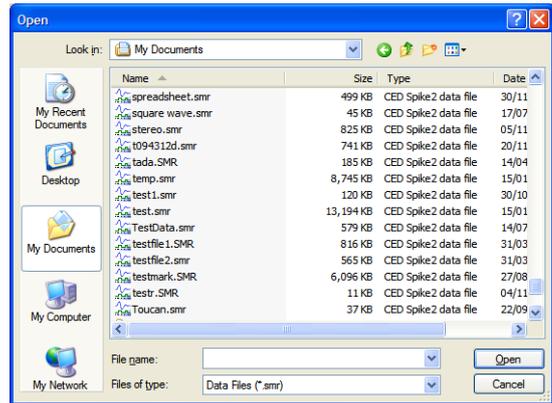
Configuration files Spike2 stores sampling configuration information in these files. The file name extension is `.s2cx`.

Filterbank files These files hold descriptions of digital filters and have the extension `.cfbx`. They are used by the Analysis menu Digital filters command.

Multimedia files These files have the extension `.avi`. They are created by the separate s2video program and are opened by the View menu Multimedia Files command.

Open

You can activate this option from the menu or toolbar. It shows the standard file open dialog for you to select a file. You can open six file types with this command: a Spike2 data file, a Spike2 result view, a text file, a script file, an output sequence file or an XY file. The type of the file is selected with the Files of type field.



When you select a data file, Spike2 looks for a resource file of the same name, but with the file extension `.s2rx`. If this is found, the new window displays the file in the same state and screen position as it was put away. Several windows may open if the file was closed with the `Ctrl` key held down. See the Close command, below, for more details.

If a text file is selected, a text window opens. This can be used as a notepad or as a repository for data copied from other parts of Spike2. If an output sequence or script or XY type is selected, an output sequence or script or XY window is created.

Import

Spike2 can translate data files from other formats into Spike2 data files. The import command leads to a standard file open dialog in which you select the file to convert. You then set the file name for the result; Spike2 will suggest the same name with the extension changed to `.smr`. The details of the conversion depend on the file type.

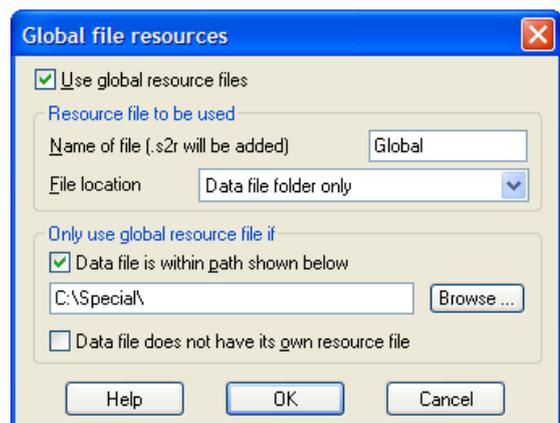
Supported formats include text files with data in columns, the CFS files used by CED programs such as Signal and SIGAVG, Spike2 for Macintosh files as well as data from many third party vendors. Spike2 searches the `import` folder in the Spike2 installation folder for CED File Converter DLLs. If you need to translate a file format that is not covered, please contact us and describe your requirements. The script language command to convert files is `FileConvert$()`.

Global Resources

If the current view is not a time, result or XY view, this command appears in the menu, otherwise it appears in the Resource Files popup menu.

Normally, each Spike2 data file has an associated resource file with the same name and `.s2rx` extension. These per-file resources remember the screen layout, cursor positions, active cursor parameters and other settings. They allow each file to open with exactly the same screen appearance it had when it closed at the cost of an extra file on disk.

However, per-file resources do not let you use the same display and cursor settings for a sequence of files. If you enable global resource files, you can use a single resource file (or one resource per folder) to control multiple data files. This is particularly powerful when you



review data stored on a read-only device such as a DVD drive. Unlike the per-file resource files, which are updated automatically whenever you close a resource file, global resource files are never updated automatically. Use **Update Global Resource** to update the current global resource and **Save Resources As** to create a new resource file with the current settings for the current view.

The use of global resources is managed from the Global file resources dialog, which has these fields:

- Use global resource files** Check this box to use global resources. If this is unchecked, no global resource files are used and each data file has its own resource file. Please remember that using this dialog makes no difference to the resources used by any open time view. It changes the resource file that is used when a saved time view opens. If you check this box and no global resource file is found, Spike2 behaves as if the box was not checked.
- Name of file** This field sets the name of the resource file. The name should not include a path or the `.s2rx` file extension. These are added automatically, as required. If this field is blank, no resource file is used.
- File location** This field sets where to search for the global resource file. You have three choices: *Spike2 installation folder only*, *Data file folder then Spike2 folder*, *Data file folder only*. The Spike2 installation folder is wherever the `sonview.exe` application file is located. The data file folder is the folder from which the data file is opened.
- Only use global resource file if** If you do not check any boxes in this area of the dialog, global resource files are applied to all data files. This dialog area allows you to restrict the files that use the global resource file in place of the per-file resources.
- Data file is within the path shown below** If you set a path and check this box, only files that lie within this path are considered for global resources. For example, if you only wanted to apply global resources to files read from your DVD drive, you might set this to `E:\` (if `E:\` is the path to your DVD).
- Data file does not have its own resource file** Check this box if you would prefer to use the per-file resources if a resource file with the correct name and in the same folder as the data file exists.
- Resource Files** This menu item replaces **Global Resources** when the current view is a time, result or XY view. It opens a pop-up menu from which you can select **Global Resources**, **Apply Resource File**, **Save Resources As** and **Update Global Resource**.
- Apply Resource File** This command applies a user-chosen resource file to the current time, result or XY view. You can use this to apply a complex window arrangement or active cursor measurement to multiple files. The script language equivalent is `FileApplyResource()`. You can also use the **Global Resources** command to automatically apply a specific resource file.
- Save Resources As** This command saves the resources associated with the current time, result or XY view to a resource file. The command opens a file save dialog for you to set the file name.
- Update Global Resource** This command is enabled for the current time, result or XY view if global resources are enabled and a global resource file name is set. It updates the global resource file with the current view settings. If the global resource file does not exist, one is created.

Close This command closes the active window. If you use this command on an unsaved window, you are prompted to save it before closing the window. However, if the **Edit** menu **Preferences** option is set to not prompt to save modified result and XY views, you will not be prompted to save these views and they will be lost when you close them.

Close all associated windows If you hold down the **Ctrl** key and close a time view, Spike2 also closes all associated windows. This does not work with a newly sampled file, save it first. If the **Edit** menu **Preferences** option is set to not prompt to save result and XY views, associated result views are deleted. However, in addition to saving the state of the data file in a `.s2r` resource file (see page 6-1), Spike2 also saves the state and contents of associated result view windows in the resource file. Next time you open the data file the result views are recreated from the resource file in the same state as they were closed.

Revert To Saved You can use this command with a text file, a script file or an output sequence file. The file changes back to the state it was in at the last save.

Save and Save As



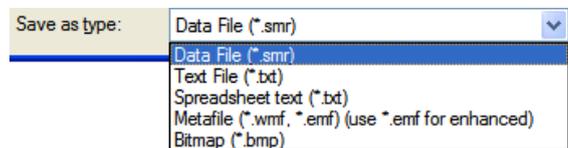
These commands are available when the current window is a text, script, output sequence, result or XY view or a time view holding newly sampled data that has not been saved. **Save** writes the file with its current name unless it is unnamed, in which case you are prompted for a name. **Save As** writes the file with a different name and gives you the option of saving the data as a different type (see **Export As** for a list of other types).

Time view files are kept on disk, not in memory, as they can be very large. Changes made to these files are permanent as they are made on disk. When you save a newly sampled data file, you give the file a name (replacing a temporary name). If you save it to a different drive from that set in the **Edit** menu **Preferences**, Spike2 copies the file to the new drive, then deletes the original. Moving a large file can take several seconds.

Sequence, text, result, XY and script files are held in memory. Changes made to these files are not permanent until the file is saved to disk.

Export As

This menu item replaces **Save As** when the current window is a saved time view. This opens a File Save As dialog. You can choose to save the file in a variety of formats: Data file (`*.smr`) to export as a new Spike2 file, Text file or Spreadsheet text (`*.txt`), Metafile (`*.wmf` or `*.emf`) for a scaleable image and Bitmap file (`*.bmp`). More types are listed if you have installed external exporters.

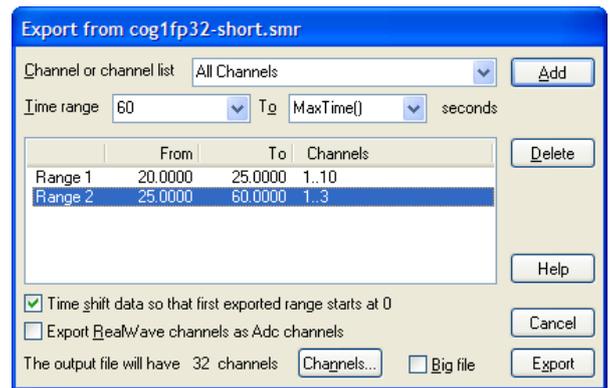


Choose a format and either select an existing file to overwrite or type a new file name, then click **Save**. What happens next depends of the export format:

External exporters

These use the same dialog as for exporting to a data file to select channels and a single time range. There may be additional dialogs depending on the target format. See the **Export** folder for additional documentation for external exporters (for example, there is documentation on the MATLAB exporter in the file `m1_export.pdf` if the MATLAB support was installed).

Data file The top section of the dialog sets channels and a time range to export. Click **Add** to append this to the list. **Delete** removes list entries. You can choose channels from the drop down list or type in a channel specification. You can export channels from the original file plus memory and virtual channels; duplicated channels are not listed. Check the **Time shift data...** box to time shift the output data so that the **Range 1 From** time becomes 0.0 in the output file.



Export RealWave channels as Adc Channels Check this to convert RealWave channels using the channel scale and offset so that pre-4.03 versions of Spike2 can read them. This is hidden for external exporters.

Channels... This sets the maximum number of channels in the output file, from 32 to 400. Spike2 versions before 4.02 only read files with 32 channels. Version 5 reads up to 256 channels (5.15 onwards can read 400). External exporters omit this field.

Big file Check this to use big file mode for the output. Only check this if the output file will be larger than 2 GB as versions of Spike2 before 6.11 cannot read big files.

Export Click this to write the selected time ranges and channels to a new data file. Channels are written in order of ascending channel number. Where possible, channels are copied to the same channel number in the output file. If this is not possible, for example for a memory channel, the channel is written to the lowest numbered unused channel in the output file.

Text file This is the same as the Edit menu **Copy As Text** command, but with the output sent to a text file and not the clipboard. See the **Copy As Text** command for details of the dialogs. This writes selected channels or all visible channels if no channel is selected.

Spreadsheet text file Copy a time view as a text file for easy import to a spreadsheet. This is the same as the Edit menu **Copy Spreadsheet** command, but with the output sent to a file. This writes selected channels or all visible channels if no channel is selected.

Bitmap, JPEG image, Portable Network Graphic, Tagged Image Format file These options copy the window to a file as bitmap at the screen resolution. If you need to scale or edit the image, a Metafile copy may be better. The Portable Network Graphic format will usually be the smallest on disk unless the screen contains a real-world image or a sonogram when JPEG may be smaller (but at the cost of artefacts around axis lines and text). Bitmap images are usually the largest.

Metafile This option copies the window as a Windows (Placeable) Metafile (*.WMF) or as an Enhanced Metafile (*.EMF). Enhanced Metafiles are theoretically better as they support cubic spline and sonogram drawing modes. However, some graphics programs do not import Enhanced Metafiles very well. The default format is a Windows Metafile. To output in enhanced format, set the file extension to .EMF. For example, saving as `fred` or `fred.wmf` saves a Windows Metafile, save as `fred.emf` for an enhanced metafile.

These file formats can be scaled without losing resolution and are the preferred format for moving Spike2 images to drawing programs. The Edit menu **Preferences...** option lets you increase (or decrease) the output resolution. This can be important when saving time view and result view data as the number of vectors produced when drawing high resolution may stop some drawing programs from reading the file.

Load and Save Configuration

These commands manage Spike2 configuration files (.s2cx extension). These files hold the full sampling setup including any output sequence file, waveform output and window arrangement and the types of on-line analysis required. The Load Configuration and Save Configuration commands transfer the Spike2 configuration between disk and the application. Both commands open a dialog to select a file for loading or saving.

You can read some sampling parameters from a Spike2 data file (.smr extension). This is useful if you have lost the corresponding configuration file. **There is no guarantee that a configuration read in this way will exactly match the data file;** you must check the result carefully and you may need to adjust the resulting configuration.

After loading a configuration, the Sampling Configuration dialog opens to display the result. If you wish to sample immediately, click the Run Now button.

You can also click a button in the Sample Bar to load a configuration and start sampling (see the *Sample menu* chapter for more information).

Default configuration files
default.s2c
last.s2c

If the configuration file `default.s2cx` is found, it is loaded when Spike2 starts. You hold down the control key while activating the File menu and use the Save Default Configuration command to create this file automatically. The standard file `last.s2cx` holds the last sampling configuration. It is saved each time sampling stops. If `default.s2cx` cannot be found, and `last.s2cx` exists, `last.s2cx` is loaded.

Page Setup

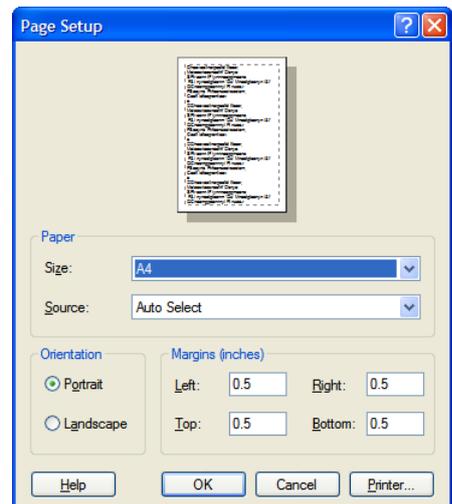
This opens the printer page setup dialog. The dialog varies between operating systems and printers. See your operating system documentation for more information. The important options that are always present include the paper orientation (portrait or landscape), the paper source (if your printer has a choice), the printer margins and the choice of printer.

The Orientation option (portrait or landscape) applies to all output except the Print Screen option, which has its own selector for the output orientation.

The printer margins will appear in inches or in millimetres, depending on the locale set for your computer. These margins are used for all printed output. The left and right margins are applied to everything, including headers and footers. The top and bottom margins apply to everything except headers and footers, which have their own top and bottom margin settings (see the Page Headers description). The top and bottom margins you set here are further modified if a header or footer would collide with them.

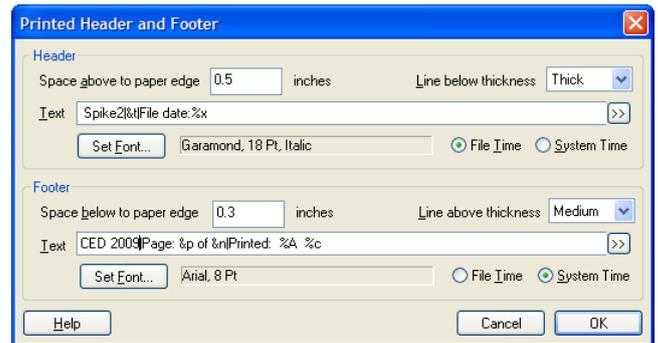
Most printers have an unprintable area near the paper edge. If you set margins smaller than the unprintable area, the margins are increased so that all your output is visible. If you set margins that reduce the printable area too much, the margins are ignored.

Registry use Spike2 saves the page margins in units of 0.01 mm in the system registry. You can find them in the `HKEY_CURRENT_USER\Software\CED\Spike2\PageSetup` folder as `REG_DWORD` values: `LeftMargin`, `RightMargin`, `TopMargin` and `BottomMargin`.



Page Headers

You can apply headers and footers to all printed output. The headers set in this dialog are used with all Time, Result, XY and text-based views. Other printed output (for example from the Print Screen command) uses all these settings except the text, which each command provides.



Header and footer positions

The horizontal position is set by the left and right margins in the Page Setup dialog. The vertical positions are set by the space above the header and the space below the footer fields in inches or millimetres, depending on the locale. If your header or footer encroaches on the top and bottom margins set in the Page Setup dialog, the top and bottom values are adjusted to keep the output clear of the header and footer.

Line thickness

You can choose between: None, Hairline, Thin, Medium and Thick. A Hairline is the thinnest line possible on the output device. The other settings should be self-explanatory. The header line is drawn below any header text, the footer line is drawn above any footer text. If there is no header or footer text, no line is drawn.

Text

This is the text to display as the header or footer. If this field is empty, the header or footer is omitted (including any line). You can split the text into left-justified, centred and right-justified sections with the vertical bar character. You can also insert codes that are replaced by document and time information. The simplest way to do this is by clicking the >> button to the right of the text and choosing an option.

Set Font

Click this button to choose a font for the header and the footer. Font sizes are limited to the range 2 to 30 points.

File/System Time

You can insert times into both the header and the footer. However, you have to choose between the current time and the file time. This allows you to display the file time in the header and the current time in the footer, or vice versa.

Document information

The following codes are replaced by document information:

&f	File and path	&F	Upper case file and path
&t	Document title	&T	Upper case document title
&p	Page number	&n	Total number of pages
&&	The ampersand sign (&)		

Time and date codes

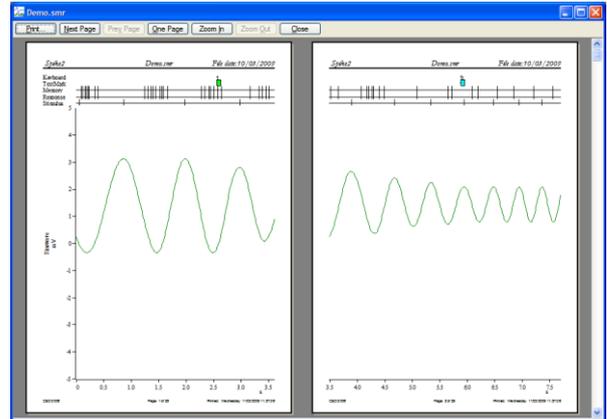
You can insert times and dates using % followed by a character code. The combination is replaced as described in the table. You can also use %#c and %#x for a long version of dates and times. You can remove leading zeros from numbers with #, for example %#j.

%a	Short day of week (Mon)	%p	Indicator for A.M or P.M.
%A	Long day of week (Monday)	%S	Seconds as number (00-59)
%b	Short month name (Jan)	%U	Week of year, Sunday based (00-53)
%B	Long month name (January)	%W	Week of year, Monday based (00-53)
%c	Date and time for locale	%w	Sunday based weekday (0-6)
%d	Day of month (01-31)	%x	Date formatted for locale
%H	Hour in 24 hour format (00-23)	%X	Time formatted for locale
%I	Hour in 12 hour format (01-12)	%y	Year without century (00-99)
%j	Day of year (001-366)	%Y	Year with century, e.g. 2004
%m	Month as number (01-12)	%z / Z	Time zone name
%M	Minute as number (00-59)	%%	The percent sign (%)

Registry use Spike2 saves the header and footer settings in the system registry. You can find them in the `HKEY_CURRENT_USER\Software\CED\Spike2\PageSetup` folder as strings: `Header`, `Header info`, `Footer` and `Footer info`. The `Header` and `Footer` items hold the text strings. The two `info` items are strings that code up the font name, point size, bold and italic settings, distance to the paper edge in units of 0.01 mm, line thickness (0=none, 1=Hairline, 2=Thin, 3=Medium, 4=Thick), and File time (0) or System time (1).

Print Preview

This option displays the current window as it would be printed by the Print option. You can preview time, result, XY and text based windows. You can zoom in and out, view a single or two pages, step through pages of multi-page documents and print the entire document using a toolbar at the top of the screen. Use the Close button to leave this mode without printing.



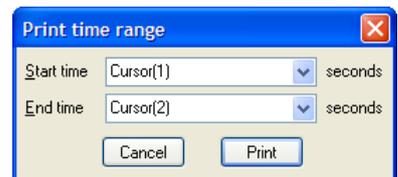
In previous version of Spike2, the preview window took over the entire program and no other commands were available apart from those on the preview toolbar. From Spike2 version 6, the preview takes place inside the frame of the time, result or XY view. You can now access the File menu to change the headers and footers and print margins. If you use the menus to make changes to the contents of the displayed data, the screen may not display the changes until you exit from Print Preview mode.

Print Visible, Print and Print Selection

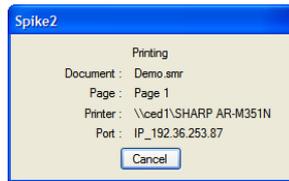


These commands print time, result, XY and text-based windows. Any scroll bar at the bottom of the window is not printed. The Edit menu Preferences dialog sets the line widths used on screen and for printing. Print Selection prints the selected area of a text window. Print Visible prints the visible data in the current window. Print prints a specified region of a time, text or result view; each printed page holds the x axis range of the window. Print in an XY view is the same as Print Visible.

To print an entire data file, set the width of the time window holding the file to the page size required in the print, then select Print. Set the start time to zero and use the drop down list to set the end time to `Maxtime()`. Beware that Print could require several hundred miles of paper to complete the printing job in the worst case! If you displayed 10 milliseconds of data across the screen in a file that is 1000 seconds long, there are 100,000 pages to print.

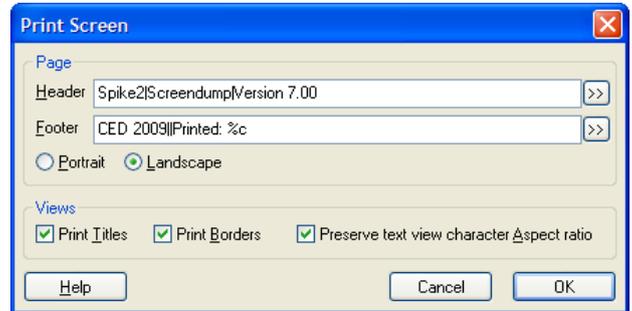


Both commands open the standard print dialog for your printer. You can set the print quality you require (in general, the better the quality, the longer the print takes) and you can also go to the Setup page for the printer. Once you have set the desired values, click the Print button to continue or the Cancel button to abandon the print operation.



During the print operation (which can take some time, particularly if you have selected a lot of data) a dialog box appears. If your output spans several pages, the dialog box indicates the number of pages, so you can gauge progress. If you decide that you didn't want to print, click the **Cancel** button.

Print Screen The Print Screen command prints all time, result, XY and text based views to one printer page. The views are scaled to occupy the same proportional positions on the printed page as they do on the screen. The page margins are those set by the Page Setup and Page Headers dialogs.



The command opens a dialog with two regions: **Page** and **Views**. In the **Page** region you can set a page header and footer and choose to print in Landscape or Portrait mode. The header text is printed with an underline across the width of the page. The footer text is printed with a line over it across the page width. The fonts used and line thicknesses are as set in the Page Headers dialog. If the header or footer is blank, both it and the associated line are omitted.

You can divide the header and footer into a left justified, centred and right justified sections with the vertical bar character, for example: `Left|Centered|Right`. You can include the current date and time in the header or footer by including, for example `%c`, as described for the Page Headers dialog. The `>>` button can be used to insert the time and date and vertical bar separators into the header and footer.

In the **Views** region you can choose to print view titles and draw a box round each view. You can also ask Spike2 to attempt to preserve the aspect ratio of characters in text windows. Without this, characters are scaled in an attempt to match the printed output to the text displayed in the view. However, this may not produce a legible output.

Registry use Spike2 saves the Print Screen settings in the system registry. You can find them in the `HKEY_CURRENT_USER\Software\CED\Spike2\PageSetup` folder. The text strings are saved as `PSHeader` and `PSFooter`. The remaining values are saved as `REG_DWORD` values of 0 (not selected) and 1(selected): `PSViewTitle`, `PSBorder`, `PSScaleText` and `PSLandscape`.

Exit This command closes all open files and exits from Spike2. If there are any text or output sequence files open that have not been saved, you will be prompted to save them before the application terminates.

Send Mail If your system has support for Mail installed (for example Microsoft Exchange), you can send documents from Spike2 to another linked computer. This option vanishes if you do not have compatible mail support.

Text-based documents and result views can be sent, even if they have not been saved to disk (Spike2 writes them to a temporary file if they have not been saved). You can send Spike2 data files, but only if they have been saved on disk.

Read-only files You can open Spike2 data files that are read-only, but you are not allowed to write any changes back to the file. Read-only data files have [Read only] added to the window title. Unless the file is on a read-only medium (such as CD) you can usually clear the read-only state by opening the Property page of the file (right click on the file name or a list of selected files) and clear the **Read-only** check box. If you copy files from a read-only drive such as a CD, the copied file is marked read-only.

Read-only text-based files open normally, but you cannot edit them.

Edit menu

The majority of the Edit menu is associated with commands that move data to and from the clipboard. You can also use these commands to search for strings in a text window or to search for the currently selected text. When the current window is a text window or an output sequence window, the Edit menu operates in the standard manner, allowing you to cut and paste text between text windows in Spike2 and other applications. When the current window holds a Spike2 data document or a result window, the behaviour is modified.

The Edit menu includes commands for searching and replacing and formatting text in text-based views. It also includes the Preferences option for the entire application.

Undo and Redo

In a text, script or output sequence window this is used to undo or redo the last text edit operation. These windows support multiple Undo and Redo operations, however you cannot undo operations that have been saved to disk or text operations that were done by a script. In Spike2 data document windows, and in result windows, you can undo most operations that change the appearance of the data window.

Cut



You can cut editable text to the clipboard in any position in Spike2 where the text pointer is visible. You cannot use this command in Spike2 data document windows or in result windows.

Copy



You can copy editable text to the clipboard plus selected fields from the Cursor Regions and Cursor Values dialogs. If you use this command in a time, result or XY window, the contents of the window, less the scroll bar, are copied to the clipboard as both a bitmap and as a metafile. See also Copy As Text.

Metafile output

To export an image to a drawing program for further annotation and manipulation, you can paste the image as either a bitmap or as a metafile. Metafiles are usually the preferred choice as you can treat the image as lines and text for further work. You can set the metafile scaling and if the image is saved as a Windows Metafile or as an enhanced Metafile in the Edit menu Preferences (see page 7-15 for more details).

Paste



You can paste text on the clipboard into a text, script or output sequence document. When you paste text, Spike2 checks that each pasted line ends with the correct end of line characters (for example, Macintosh and Windows use different sequences). The paste operation corrects incorrect sequences. To correct text that has come from a different operating system and that looks strange, select all the text, cut, then paste.

Delete

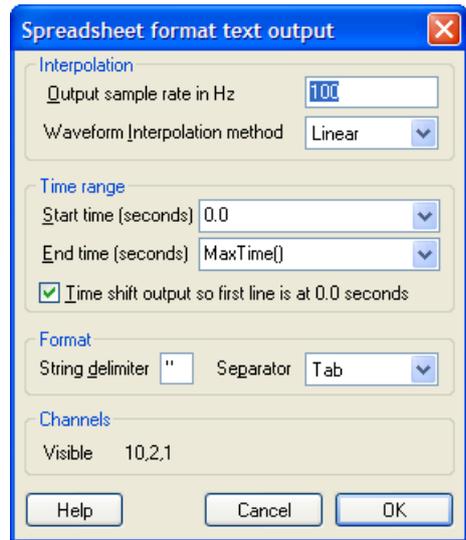
This command is used to delete the current text selection, or if there is no selection, it deletes the character to the right of the text caret. Do not confuse this with Clear, which in a text field is the equivalent of Select All followed by Delete.

Clear

When you are working with editable text, this command will delete it all. If you are in a result window, this command will set all the bins to zero and redraw the window contents. Clear removes everything; Delete removes the current selection.

Copy Spreadsheet

This copies selected time view data channels to the clipboard as text. Use **Export As** to copy to a text file. If there are no selected channels, it copies visible channels. Data is written in columns. The first column holds the time of each row (in seconds). The remaining columns hold the data, one column per channel. The first line holds titles to identify the data.



Channel output

Where possible, the output matches the display. Channels with a y axis output values that match the displayed data. Channels drawn in State mode output the state code as text. Level event channels drawn as levels output the level as a 1 or a 0. Sonogram draw mode is output as waveform mode. All other drawing modes output the number of events that fall from the time of the current line up to the time of the next line.

Interpolation

Spike2 allows each channel to have independent rates and types. To make Spike2 data easily accessible to other programs we resample the data to a common rate, set by the Output sample rate in Hz field.

In general, the sample rates of Waveform, RealWave and WaveMark data will not match the output rate you have chosen and the program must interpolate. You can choose between Linear, Cubic spline interpolation, or use the value at the Nearest data point. Waveform data drawn in Sonogram mode is treated as a waveform. All other channels that need interpolation use linear interpolation.

Time range

The Start time and End time fields set the time range to copy. You can also shift the data so that the first line has a time of 0.0 to make comparisons of time ranges easier.

Format

All values are written in columns. Each column is either numeric or text. Between each column there is a separator; choose from Tab, space or comma in the Separator drop down list. The String delimiter character is placed around text output.

Copy As Text (Result view)

This command is available in result views. It copies the bin values *in the window* for all visible channels to the clipboard as text. To copy all bins, double click the x axis and select **Show All**, then copy. The first output column holds the x value at the left of each bin. Each channel contributes one column of bin contents plus a column of error bar sizes if error bars are enabled and a column of bin counts if bin counts are enabled. The first line holds column titles. This command does not copy channels drawn as raster data. The first example is from a result view holding two waveform averages:

```
"Time"      "Filtered"    "Sinewave"
0           0.84302088   3.2980477
0.01       0.27032173   2.2613753
0.02       -0.45484864  1.043198
```

The columns are separate by Tab characters. The second example is from the same data with error bars enabled and drawn in SEM mode.

```
"Time"      "Filtered"    "SEM"        "Sinewave"    "SEM"
0           0.84302088   0.038987886  3.2980477     0.032728175
0.01       0.27032173   0.021425654  2.2613753     0.036360926
0.02       -0.45484864  0.027970654  1.043198      0.056437311
0.03       -0.74969506  0.037386934  0.63429488    0.050863126
```

**Copy As Text
(XY view)**

This menu item is available in XY views. It copies the visible points for visible channels to the clipboard. If the view has one channel or the Measurement system created it with the All channels use same X option, the output is a rectangular table with the first line holding column titles. There is one column of x values followed by one column per channel for each y value. The columns are separated by a tab character:

```
"X"      "Channel 1"  "Channel 2"
0        0.244140625  0.0170616301
2        3.122558594  0.0001053187043
4        2.211914063  0.0006319122258
6        1.889648438  -0.0005265935215
```

In all other cases, channels are output separately. For each channel, the first output line holds "Channel : cc : nn" where cc is the channel number and nn is the number of data points. The data points are output, one per line as the x value followed by the x value, separated by a tab character.

```
Channel : 1 : 5
0.0170616301      0.244140625
0.0001053187043  3.122558594
0.0006319122258  2.211914063
-0.0005265935215 1.889648438
Channel : 2 : 5
0                  0.0170616301
2                  0.0001053187043
4                  0.0006319122258
6                  -0.0005265935215
```

**Copy As Text...
(Time view)**

In a time view the command opens a dialog in which you specify the time range of data to copy, how you want the data to be copied and the output format. Select the channels to copy in the time view before you use this option. If there are no selected channels, Spike2 copies visible channels. Use the File menu Export As command to write the output to a text file.

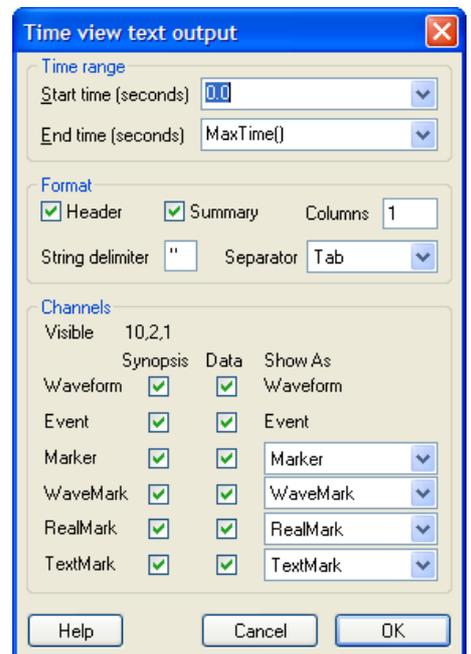
You can enable and disable a header section for the entire output and a summary section of all the channel information. You can also set the number of columns to be used when writing waveform and event times and the separators used between fields and to delimit text. For each type of channel you can enable the channel synopsis and data output. You can also choose to dump a channel type in its native format, or in any compatible format. Output sections are preceded by a keyword so that other programs can parse the file.

Time range

You select the time range of data to copy as text in this dialog. You can either type in the range, or use the drop down lists to select start and end times for the data output. When you are satisfied with the range, click on Cancel or OK to start the output.

Format

All information is written in fields. Each field is either numeric or text. Between each field there is a separator, which can be set to be a Tab character, a blank or a comma in



the **Separator** drop down list. Most programs that accept tabulated numbers will accept space, a Tab or a comma. The examples below use space as a separator.

A numeric field holds numbers only, either floating point numbers with a decimal point, or integer numbers. A text field is a sequence of characters that may include spaces. Text fields may hold numbers, but numeric fields cannot hold text. You can mark the start and end of a text field with a special character (usually ") so that a program reading the field can include blanks (spaces) and punctuation within a field without confusion. All keywords are treated as text fields. You can set a one character delimiter in the **String delimiter** field, or leave it blank for none. The examples use " as a delimiter.

Columns You can choose to write the data for waveform channels and event times in columns. This does not write one channel per column (use Copy Spreadsheet for that format).

Header The first part of the output is a header that displays information about the data file. This is given after the keyword **INFORMATION**. The header consists of the file name followed by five lines of the file comment and a blank line. The header block is always 8 lines long (if present); blank file comment lines are not skipped.

```
"INFORMATION"
"demo.smr"
"These five lines"
"contain"
"the file comment"
"for this"
"data"
```

Summary The summary section starts with the keyword **SUMMARY** followed by a summary of the channel information for each selected channel. The information given for the channel varies with the channel type. The first three fields are the same for all channels, being the channel number, the channel type and the channel title. The remaining fields are:

```
Waveform    Units, Ideal rate, Actual rate, Scale, Offset
WaveMark    Units, Ideal rate, Actual rate, Scale, Offset, Points, PreTrig
Event        Predicted mean rate
Marker       No other information
```

Here is some typical output. The **SUMMARY** section is terminated by a blank line.

```
"SUMMARY"
"1" "Waveform" "Sinusoid" "Volts" 100.0 100.0 1.0 0.0
"2" "Evt-" "Synch" 5.0
"4" "WaveMark" "Shapes" "Volts" 100.0 100.0 1.0 0.0 30 5
"31" "Marker" "untitled"
```

Channel information

This section lists the channels that can be output and lets you choose the information to dump for each channel. You can also choose how much data is output for marker-derived data types. For each channel the output always contains an information section, headed by the keyword **CHANNEL**, followed by the channel number. The channel number is a text field. The next three lines are output if the **Synopsis** box is checked and contain the channel data type, comment and title. The information section ends with a blank line.

```
"CHANNEL" "1"
"Waveform"
"Signal generator"
"Sinusoid"
```

The information that follows varies with the channel type:

Waveform A channel containing waveform data has the channel information followed by the data. The data starts with the channel units and ideal sample rate. The next line contains the keyword **START** followed by the start time of the data in seconds and the time increment

per data point, also in seconds. It is possible for gaps to occur in the waveform data. A gap is shown by a line with the word `GAP` followed by the start time of the new section and the time increment. The data then follows as a list of waveform values.

```
"CHANNEL" "1"
"Waveform"
"Signal generator"
"Sinusoid"

"Volts" 100.0000
"START" 0.00 0.01
-3.7208
-3.8356

"GAP" 4.23 0.01
-2.2901
-2.6075
-2.8858
```

Event Event data comes in three types, rising edge, falling edge and both edge triggered (`Evt-`, `Evt+` and `Level`). These three types are represented in basically the same way.

```
"CHANNEL" "2"
"Evt-"
"Synchronisation pulse from Signal generator"
"Synch"

0.2022
.
.
```

In the case of an event channel triggered on both edges, the state before the first transition is shown at the start of the data, by one of the words `HIGH` or `LOW`. `HIGH` means that the first time in the list represents a low to high transition, `LOW` means the reverse. If there are no times in the list, these labels are the opposite of the state in the period.

```
"HIGH"
0.2022
.
.
```

Marker Marker data consists of a series of times and 4 bytes of marker information. The channel synopsis is as for other channels, and the data is displayed as a time followed by a string of 4 characters that represent the 4 bytes of information. If the bytes are non-printing characters a `?` is shown instead. The 4 bytes are also displayed as numbers after the string. You can force marker data to be dumped as event data.

```
"CHANNEL" "16"
"Marker"
"No comment"
"untitled"

1.0906 "p???" 112 0 0 0
3.0336 "a???" 97 0 0 0
5.9802 "u???" 117 0 0 0
6.9018 "l???" 108 0 0 0
```

WaveMark WaveMark data is structured as a marker plus a short section of waveform data. The channel synopsis is the same as for waveform data, but also has the number of waveform points associated with each event and the number of pre-trigger points. The dump of data starts with a line holding the channel units, the ideal sampling rate, the number of data points and the number of pre-trigger points:

Units, Ideal rate, Points, PreTrig

This is followed by blocks of data, one for each WaveMark event in the time range. The first line of the block holds the keyword WAVMK, followed by the time of the first data point in the event, the time interval between waveform points and the four marker bytes:

```
WAVMK,eventTime,adcInt,mark0,mark1,mark2,mark3
```

The following lines in the data block hold the waveform data values. This block is repeated for each event in the time range. A typical WaveMark channel begins:

```
"CHANNEL" "1"  
"WaveMark"  
"Channel comment for demonstration"  
"Title"
```

```
"Volts" 100.0 30 5  
WAVMK 0.0210 0.01 1 0 0 0  
0.1318  
.  
.  
"WAVMK" 0.8110 0.01 3 0 0 0  
.  
.
```

You can also force a WaveMark channel to be dumped as though it was a marker channel, as an event channel or as a waveform channel.

RealMark The first line of RealMark data output holds the channel units followed by the number of reals attached to each event. This is followed by data blocks flagged by REALMARK followed by the event time and the four marker bytes. The following lines in the data block hold the real number. You can force RealMark data to be dumped as Marker or event data. Typical data follows:

```
"Units" 2  
"REALMARK" 1.0906 112 0 0 0  
1.8923  
4.8567  
  
"REALMARK" 1.8603 113 0 0 0  
1.8224  
4.123
```

TextMark The first line holds the maximum number of characters allowed. This is followed by data blocks separated by blank lines, one for each TextMark. The first line of each block starts with TEXTMARK, followed by the marker time and the four marker bytes. The second holds the text. You can force TextMark data to be dumped as Marker or Event data.

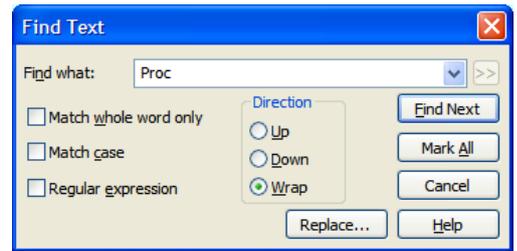
```
100  
"TEXTMARK" 0.5234 112 0 0 0  
"This is where I added the secret ingredient"  
  
"TEXTMARK" 9.8603 113 0 0 0  
"Control point 1"
```

You can also export this data type by double clicking any TextMark and copying the items from the TextMark dialog.

Select All This command is available in all text-based windows and selects all the text, usually in preparation for a copy to the Clipboard command.

Find Find Again Find Last

The Edit menu Find command opens the Find Text dialog. The dialog is shared between all text-based views. It is closely linked to the Find and Replace Text dialog and shares all its fields with it; opening this dialog closes the Find and Replace dialog. Click **Replace...** to swap to the Find and Replace dialog. A successful search moves the text selection to the next matching string. Searches are line by line; you cannot search for text that spans more than one line.



Find Next starts a search for the text in the **Find what** field. The **Mark All** button sets a bookmark on all lines that match the search text, but does not move the selection. Searches are insensitive to the case of characters, unless you check **Match Case**. Check the **Match whole word only** box to restrict the search so that the first search character must be the start of a word and the last search character must be the end of a word.

Search direction Select **Up** to search backwards, towards the start of the text. Select **Down** to search forwards towards the end of the text. Select **Wrap** to search forwards to the end, then wrap around to the start and stop when you reach the current position. Searches do not include the currently selected text.

Regular expression Check this box to search for *regular expressions*. This disables **Match whole word only** as regular expressions have their own way to match word starts and word ends. It also disables **Up** searches; regular expression searches are forwards only. It enables the **>>** button, which displays a list of regular expressions to insert into the search string. The simplest pattern matching characters are:

- ^ Start-of-line marker. Must be at the start of the search text or it just matches itself. The following search text will only be matched if it is found at the start of a line.
- § End-of-line marker. Must be at the end of the search text or it just matches itself. The preceding text will only be matched if it is found at the end of a line.
- .

To treat these special characters as normal characters with regular expressions enabled, put a backslash before them. A search for “`^\^.\.`” would find all lines with a “`^`” as the first character, anything as a second character and a period as the third character.

You can use `\a`, `\b`, `\f`, `\n`, `\r`, `\t` and `\v` to match the ASCII characters BELL, BS (backspace), FF (form feed), LF (line feed), CR (carriage return), TAB and VT (vertical tab). Searching for `\n` and `\r` will not normally match anything as `\n` or `\r\n` mark line ends and the search is of complete lines ignoring end of line markers. `\xnn` can be used to search for an ASCII character with hexadecimal code `nn`.

To search for one of a list of characters, enclose the list in square brackets, for example `[aeiou]` finds a vowel. A hyphen defines a character range: `[a-zA-Z0-9]` matches any alphanumeric character. To include the `-` character in a search, place it first or last. To include `]` in the list, place it first. To search for any character that is not in a list, place a `^` as the first character. For example, `[^aeiou]` finds any non-vowel character.

You can search for the start and end of a word with `\<` and `\>`. Word characters are the set `[a-zA-Z0-9]`, or `[a-zA-Z0-9%$]` in script views. You can also use `\w` to match a word character and `\W` to match not a non-word character. Likewise, `\d` matches a decimal number and `\D` to matches a non-number character and `\s` matches white space (space, TAB, FF, LF and VT) and `\S` matches non-white space. You can use `\w`, `\W`, `\d`, `\D`, `\s` and `\S` both inside and outside square brackets.

Regular expression multiple matches There are search characters that control how many times to find a particular character. These characters follow the character to search for:

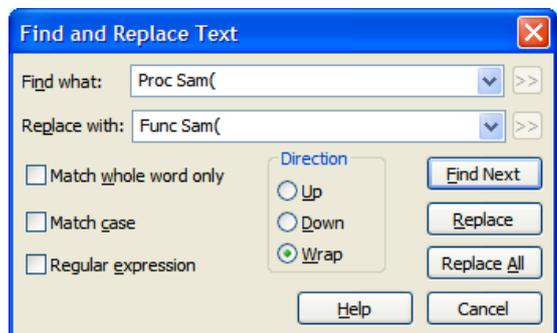
- * Match 0 or more of the previous character. So `51*2` matches `52`, `512`, `5112`, `51112` and `h.*1` matches `h1`, `hel`, `hail` and `B[aeiou]*r` matches `Br`, `Bear` and `Beer`.
- + Match 1 or more. The same as “*”, but there must be at least one matching character.
- ? Match 0 or 1 of previous character. So `51?2` matches `512` and `52`.

To use these as normal characters for matching, put a backslash before them.

Regular expression tagging You can tag sections of matched text by wrapping them in `\(` and `\)`. You can then insert the tagged text later in the regular expression (or in the replace text in the Find and Replace dialog) using `\n`, where `n` is 1 for the first tag, up to 9 for the ninth. For example, `\(foo\) - \1` matches `foo-foo`. More interestingly, `\(\<[a-zA-Z]+\>\) - \1` matches `Jim-Jim`, `plum-plum` and the like.

Find Again, Find Last Find Again and Find Last repeat the current search forwards or backwards.

Replace The Edit menu Replace command is available when the current view is text-based. It opens the Find and Replace Text dialog in which you can search for text matching a pattern and optionally replace it. The search part of the dialog is identical to the Find Text dialog. The search pattern set by the Find what field can be a simple match, or can be a regular expression. In regular expression searches, the replacement text can refer back to tagged matches in the search text. See the Find Dialog for details of regular expressions and tagged matches. The >> buttons are also enabled in regular expression mode and let you insert expressions into the search and replace text.



Replace with This field holds the text to replace the matched search text. In a regular expression search, you can include tagged matches from the search text using `\1` to `\9` as described for the Find dialog. For example, suppose you have variables named `fred0` to `fred17` that you want to convert into an array `fred[0]` to `fred[17]`. You can do this by setting the Find what field to `\<fred\ (\d+)\>` and the Replace with field to `fred[\1]`.

Replace The Replace button tests if the current selection matches the Find what field and if it does, the field is replaced by the Replace with text field and the selection is moved to the next match. If the selection does not match, Replace is equivalent to Find Next.

Replace All This button searches for all matches in a forward direction starting from the beginning of the text to the end, and replaces them.

Edit toolbar The Edit toolbar gives you access to the edit window bookmarks and short cuts to the Find, Find next, Find previous and Replace commands. If you are unsure of the action of a button, move the mouse pointer over the button and leave it for a second or so; a “ToolTip” will reveal the button function and any short-cut key associated with it.



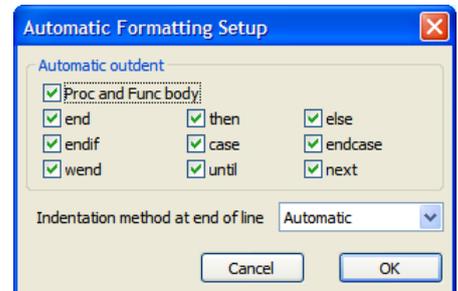
The four bookmark functions toggle a bookmark on the current line, go to the next or previous bookmark and clear all bookmarks. You can set bookmarks on all lines containing a search string with the Edit menu Find command.

If the Edit toolbar is not visible, click the right mouse button on an empty area of a toolbar or of the Spike2 main window. Then select Edit bar in the pop-up menu. You can use the same procedure to hide it. The bar can be docked to any side of the application main window or dragged off the application main window as a floating bar.

Auto Format Automatic formatting is available for script views and applies standard formatting while you type and lets you reformat entire scripts or selected script regions. Formatting is done by indenting lines of text based on the script keywords. An indented line is one that starts with white space. The indenting is in units of the Tab size set for the view in the Script Editor Settings dialog. Indentation is done with Tab characters if you have chosen to keep Tabs in the view, otherwise indentation is done with space characters. There are two sub-commands:

Apply Formatting If any text is selected in the current view, all lines included in the selection are formatted. If there is no selection, the entire document is formatted. If text is selected, you can right-click in the text view and choose Auto Format Selection from the pop-up menu to activate this option.

Settings... The Automatic Formatting Setup dialog controls how code is formatted. Formatting is based on the same scheme that is used for folding text. Each script keyword that starts a block construction (proc, func, if, docase, while, repeat, for) increases the indent, and the keywords that complete blocks decrease the indent. However, many users prefer the look of the text with extra outdents.



CED standard formatting has all the boxes checked, however, it makes no difference to the script operation; what you choose is a matter of personal taste. It is a good idea to use consistent indenting as it helps you to understand the structure of the script.

Proc and Func body If this box is not checked, all text within a function or procedure is indented. Check the box to outdent the text between the Proc or Func and the end.

end...next You can choose to outdent any line starting with one of the keywords end, then, else, endif, case, endcase, wend, until and next.

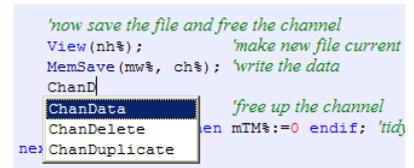
Indentation method at end of line This field determines what happens to the indentation of the current line and the new line when you type the Enter key. The settings are:

- None No automatic formatting is applied. The new line is not indented.
- Maintain The new line is indented to match the indentation of the current line.
- Automatic The indentation of both the current line and the new line is adjusted based on the Auto Format settings.

Toggle Comments This is available in script and text sequencer views, and also in the context menu when there is a selection. It adds or removes a comment marker at the start of the line for all lines in the current selection based on the first character of the first selected line.

Auto Complete

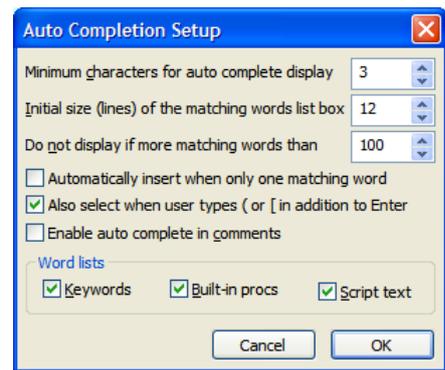
In any script, you will find that you are typing the same text items repeatedly. The editor can save you some time by popping up lists of known words that match your typing. The matching is done by looking for a word break in the text before the text caret, then matching your typing against various categories of words known to the script. Matched words are displayed in alphabetical order and the current word is highlighted.



You can use the up and down arrow keys to choose an item in the list and the Enter key to select an item or double click an item to enter it. The Esc key cancels the list. Alternatively, you can just keep on typing, which will narrow the choice of words to match. The Edit menu Auto Complete Setup dialog gives you some control over the words that are matched and when the auto-completion lists appears.

The Word lists section of the dialog lets you choose the categories of words to match your typing against. If you do not want to display auto-completion lists, clear all the categories.

You can choose from script keywords (Proc, Func, EndCase, ...), built in functions and procedures (SampleSequencer, NextTime, ...), and words that already exist in the script. You may wish to disable the Script text option if you are working with a huge file and you notice an appreciable delay after typing before the list appears.



Minimum characters for auto complete display

This sets the number of characters in a name that must be typed before the list will appear. You can set this to 1 to 12 characters. Setting a low value can cause the pop-up display to become annoying. You need to choose a count that gives you enough help, but not too much. Try a value of 3 to start with.

Initial size (lines) of the matching words list box

This field sets the maximum number of words to display at a time in the range 1 to 40. If there are more matching words, the list contains a scroll bar. After the list appears, you can re-size it by clicking and dragging on the horizontal edge furthest away from the matched text.

Do not display if more matching words than

If there are more matching words than you set here, the list is not displayed. You can set this value in the range 1 to 200 words.

Automatically insert when only one matching word

Because of the design of the script language, apart from variable declarations, all words you type in are likely to have been defined already. If you set this option, once your typing has reduced the list size to 1, the word is automatically inserted. You may want to increase the Minimum characters for auto complete display if you enable this option.

Also select when user types (or [

Normally, the list text is inserted when you type the Enter key or double-click the list text. If you check this option, the text is also inserted when you type an opening round or square bracket, followed by the bracket.

Enable auto complete in comments

Normally, automatic word completion is disabled in a comment. However, if refer to script functions in your code documentation you may find it useful to check this box.

Preferences

The Edit menu Preferences dialog has tabs for general preferences, display options, sampling options, signal conditioner setup, Spike2 time scheduling and for compatibility with previous versions of Spike2. The preferences are stored in the Windows registry and are user specific. If you have several different logons set for your computer, each logon has its own preferences. It is possible to change the preferences from a script; see the `Profile()` command for details.

General

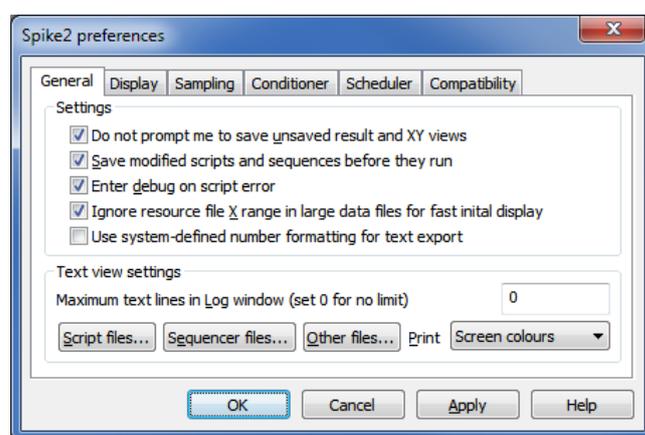
This tab holds editor preferences and general settings for saving modified views and to control what happens when an error is found in a running script.

Do not prompt me to save unsaved result and XY views

As it is often possible to recreate result and XY from the raw data, you can check this box to suppress the normal Windows behaviour of prompting you to save unsaved data.

Save modified scripts and sequences before they run

If you check this box and run a modified script or sequence, it will be saved first. If a file has no name, the File Save dialog opens to prompt for one.



Enter debug on script error

Normally, if a running script has a problem, it stops and an error message is displayed in the script window and in the Log view. If you check this box, the script debugger is activated on a script error and you can inspect the local and global errors and the call stack at the time of the error. You are not allowed to continue running the script.

Ignore resource file X range in large data files for fast initial display

Normally, when Spike2 opens a data file with an associated resource file, the same x range is displayed as was visible when the file closed. With huge files this initial display can take several seconds. Check this option to display only the first second of data if a file is a significant size.

Use system-defined number formatting for text export

Check this to use the number formatting set in your operating system when exporting text (to a text file or the clipboard). This allows the use of a comma as a decimal point.

Text view settings

This area of the dialog holds controls that let you configure the appearance of text-based views (script, text sequencer, log views and text views created by a script) on screen and when printing. The dialogs accessed by the **Script**, **Sequencer** and **Other** files buttons will apply any changes globally (to all open files of the selected type and to all future files). The same dialogs can be used from the View menu **Font** commands to make changes to the current view (there is then an extra button to apply changes to all views).

Print

This sets how to print coloured text from a text view. The choices are:

Screen colours Use the displayed screen colours. This can be wasteful of ink or toner if the background is not white.

Invert light If you display your text as a light colour on a dark background, this setting prints on a white background and inverts the text colours.

Black on White Prints black text on a white background.

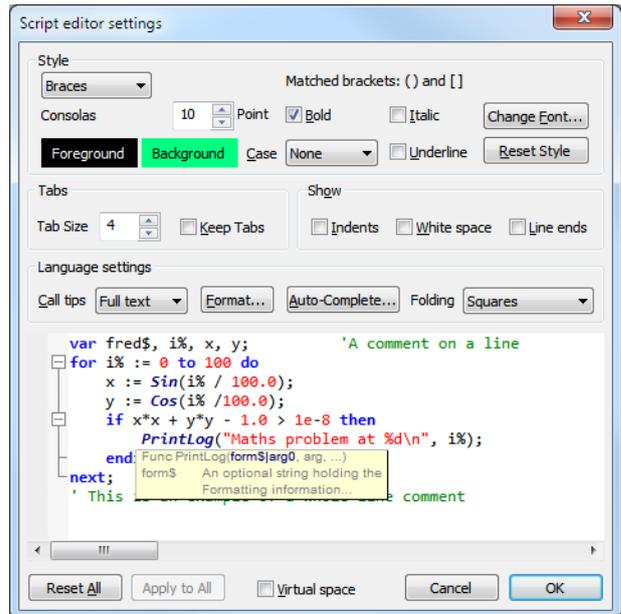
Colour on White Prints in screen colours on a white background.

Maximum text lines in Log window

You can use this field to prevent huge amounts of text accumulating in the Log view. If there are more lines than set here in the Log view and a script writes more, the oldest lines are deleted to leave this number. Set 0 for no line number limit.

Script files... Open the editor settings dialog from Edit menu Preferences General tab to change settings for all views of the current type, or from the View menu Font command to change only the current view. When used for the current view, an extra button Apply to All appears at the bottom of the dialog to apply changes to all views. The Reset All button returns the dialog values to standard settings.

The bottom half of the dialog holds example text suitable for the view type to show the effect of any changes.



Virtual space If you enable virtual space, you can move the text caret beyond the end of a line. If you add characters in virtual space, the editor will add spaces up to the added character.

Style Each view type supports one or more text styles. For each style you can set a font (including proportionally spaced fonts). The font includes the size in points (in the range 2-256) and bold and italic settings. You can also choose to force upper or lower case and underlines for all text in the style. You can set a foreground and background colour for the style by clicking on the Foreground and Background rectangles.

All views have the Default style that the remaining styles are based on; it is used for everything that is not covered by one of the other styles including the caret colour. The Tab size is based on the width of a space character in this style. If you change any aspect of this style, all the other styles that match that aspect will also change.

In a Script view, you can control the appearance of many different aspects of the display, based on the syntax of the language. There are settings for:

- White space** Style used when drawing spaces and tabs and control characters.
- Comment** The style used when displaying comments.
- Numbers** The style to use when displaying numbers.
- Keywords** The style to use for script keywords, such as `for`, `next` and `end`.
- String** The style for literal strings, such as `"This is a string"`.
- Function** Used for built-in script functions, such as `PrintLog()`.
- Operator** The style for script language operators, such as `+`, `-` and `+=`.
- Identifier** The style for function, procedure and variable names created in a script.
- Line numbers** The style used for line numbers and the gutter. When you select this item, line numbers will appear in the example text area.
- Call tips** The style for pop-up call tips. This style has an extra Tip highlight colour field, used to highlight the current argument in a function. This replaces the Bold, Italic, Case and Underline fields, which are not displayed.
- Braces** The style used to highlight matching braces, which are `()` and `[]`. Braces are highlighted when the text cursor is next to a brace.
- Bad braces** The style used to highlight an unmatched brace.

The Reset Style button reverts the current style to standard settings.

Tabs This dialog region controls the size of tabs (set in units of the width of a space character in the **Default** style) and if Tabs are implemented by saving a Tab character in the text (check **Keep Tabs**) or are implemented as multiple spaces (clear the **Keep Tabs** checkbox). If you use a fixed pitch font, such as `Courier New`, then it does not matter too much if you choose to keep Tab characters or not. If you use a proportional font for anything except comments, it is better to keep the Tab characters. When automatically formatting a script, the **Keep Tabs** setting determines if indents are generated with Tab characters or spaces.

Show In addition to displaying the text, you can also choose to display information about the white space in your file. The settings are:

Indents It can be useful when manually lining up indented loops in a script to see the indent level. Check this box to display vertical lines at each tap stop in the leading white space of each line.

White space Check this box to display spaces with a centred dot and tabs as right arrows.

Line ends Check this box to see the Carriage Return (CR) and Line Feed (LF) characters that mark the end of a line. This can be helpful to understand what is happening when working with scripts that move the caret around.

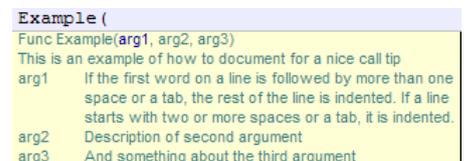
Language settings This area of the dialog is used by script and output sequencer views. It sets your preferences for call tips, automatic formatting, automatic text completion and folding. The **Format** and **Auto Complete** buttons duplicate Edit menu commands, and are described separately.

Call tips A Call tip is a block of text that pops up in a script view when you type the opening brace of a function or procedure name. The call tip text contains a synopsis of the command and a list of the command arguments, with the current argument highlighted. You can choose from **None** (no call tips), **Single line** (just the command name and arguments) and **Full text** (includes a command synopsis).

For built-in functions, the call tip text holds the command name and arguments plus one sentence of description. For user-defined `Proc` and `Func` items, the arguments are taken from the line containing the `Proc` or `Func` keyword (expected to be the first item on the line). If the line before the `Proc` or `Func` is a comment, up to 10 lines of comment are used as a description. If the line before is blank, and the line after is a comment, up to 10 lines following are used as a description. If you document your functions and procedures as in this example, they will generate tidy call tips (and fold away to a single line):

```
Func Example(arg1, arg2, arg3)
'This is an example of how to document for a nice call tip
'arg1 If the first word on a line is followed by more than one
'     space or a tab, the rest of the line is indented. If a line
'     starts with two or more spaces or a tab, it is indented.
'arg2 Description of second argument
'arg3 And something about the third argument
var x,y; 'the start of the code
...
```

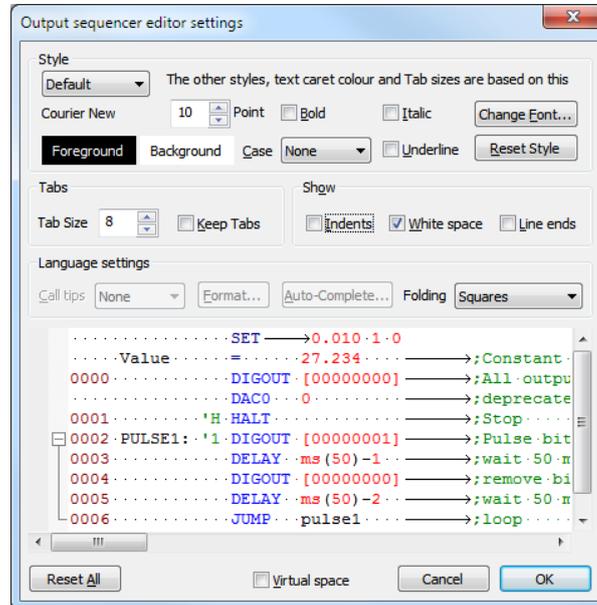
The example text above would produce a call tip looking like this. The comment markers at the start of each line are removed, and the multiple spaces after the first word or at the start of each line are replaced with a Tab character.



Folding Script views and output sequence views can display a folding margin, and allow you to fold the code by clicking in the margin, from the **View** menu **Folding** command, and by

right clicking in the view and selecting **Toggle all folds**. In a script view, fold points are based on a lexical analysis of the script text. You can choose from one of four folding styles, or have no folding margin.

Sequencer files...



The editor settings dialog for output sequencer files is very similar to that for script files. The example text in the lower half of the dialog displays some typical sequencer code, and there is no support for Call tips, auto formatting or automatic completion.

There is a list of styles that you can apply to the different elements of the sequencer text. The **Default** style is used in exactly the same way as for script views as the basis for all other styles and as the font that is measured to set the Tab size. The remaining styles are:

White space	Style used when drawing spaces and tabs and control characters.
Comment	The style used when displaying comments.
Numbers	The style to use when displaying numbers and digital i/o expressions such as [...0101].
Keywords	The style to use for standard output sequencer instructions.
Deprecated	The style to use for outmoded output sequencer instructions like <code>DACO</code> that have been replaced and may not be supported in future revisions.
Display	Text introduced by <code>></code> for display on screen during sampling.
Directives	Items such as <code>SET</code> and <code>VAR</code> that do not generate output instructions.
Operator	The style for mathematical operators like <code>+</code> , <code>-</code> , <code>*</code> and <code>/</code> .
Identifier	The style for user-defined variable names and labels.
Step	The style for the optional step numbers at the start of an instruction line.
Key	The style for keyboard links introduced by a single quote, for example <code>'H</code>
Functions	The style for built-in conversion functions like <code>ms()</code> .
Line numbers	The style used for line numbers and the gutter. When you select this item, line numbers will appear in the example text area.
Braces	The style used to highlight matching braces, which are <code>()</code> and <code>[]</code> . Braces are highlighted when the text cursor is next to a brace.
Bad braces	The style used to highlight an unmatched brace.

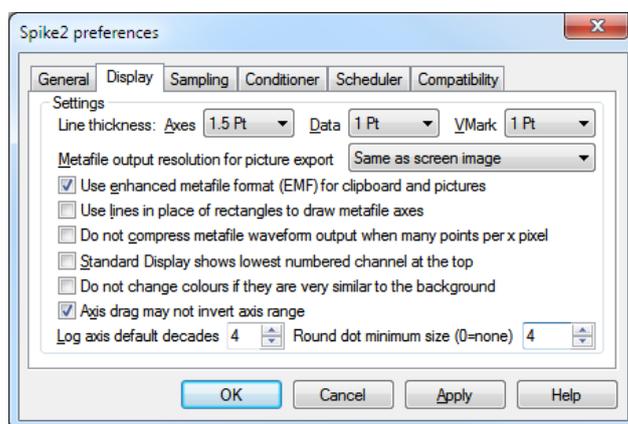
Folding support The output sequencer editor supports code folding based on the keyboard link entry points. That is, you can fold up all code from a line holding a keyboard link up to the next line holding a link.

Other files... The editor settings dialog for all views except script and sequencer views is the same as the dialog for script views, except that there are no language settings and there are only the Default and Line numbers text styles.

Display This dialog tab contains options for data display, printing and image export.

Line thickness: Axes, Data and VMark (vertical markers)

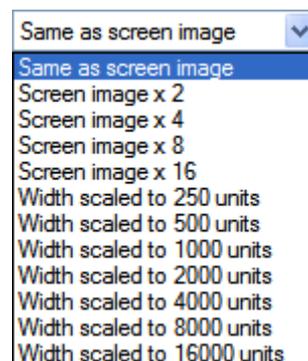
Choose from Hairline (as thin as possible) and a list of point sizes. A point is 1/72nd of an inch, about one screen pixel. If coloured lines do not print on a monochrome printer, increase the line thickness or use the View menu Use Black and White command. Click



Apply to redraw screen images to show the effect of any change. WARNING: wide lines take longer to draw than a single pixel line. You can override the pen widths of individual channels with the View menu Channel Pen Width command. The draw mode dialog sets dot sizes as multiples of the Line thickness you set here for data.

Metafile output resolution

Spike2 saves time, result or XY views as pictures in either bitmap (a copy of the screen image) or metafile format. A metafile describes an image in terms of drawing operations. You can set the metafile drawing resolution; the higher the resolution, the more detail in the picture.



A problem for time and result views with many data points is that the higher the resolution, the more lines need to be drawn, and many drawing programs have limits on the number of lines they can cope with. You can use multiples of the screen resolution, or fixed resolutions. If you are not sure what setting to use, start with *Same as screen image* and adjust it as seems appropriate.

Use enhanced Metafile format

Spike2 supports two metafile formats: Windows Metafile (WMF) and Enhanced Metafile (EMF). WMF is a relic of 16-bit Windows and has limitations, but is widely supported. EMF is the standard for 32-bit programs and has many more features. However, some graphics packages do not support this fully.

If you use EMF format you can export waveforms drawn in cubic spline mode and sonogram data as part of the metafile. If you use the WMF format, waveforms drawn in cubic spline mode will be exported as plain lines and sonograms will be blank.

Use lines in place of rectangles...

This affects metafile output. Some graphics programs cannot cope with axes drawn as rectangles; check this box to draw axes as lines. We use rectangles to make sure that axes drawn with pens of other than hairline thickness join neatly.

Do not compress metafile waveform output

When Spike2 draws waveform data on the screen, it does not waste time with lines that make no visual difference to the output. If there are more than 3 data points per horizontal pixel, Spike2 draws one vertical line per horizontal pixel to give the same visual effect as all the separate lines. However, when you export an image as a metafile

this may not look correct as it relies on a perfect match between the vertical line width and the vertical line separation.

If you check this box, no compression is done when data is saved as a metafile. For the most precise image, set the Metafile output resolution to 16000 units and check this box. When you import a metafile, most drawing programs will preserve the waveform as connected lines. Spike2 writes long sequences of data points in blocks of up to 4000 points. The last point of a block is at the same position as the first point of the next.

Standard display shows the lowest numbered channels at the top

You can change the order of time and result view channels by clicking and dragging channel numbers. This sets the order when you use the Standard display command or open a new window. If the box is unchecked, lower numbered channels are at the bottom.

Do not change colours if they are very similar to the background

Spike2 checks how similar the colours of items such as channel data and text are to the background. If they are too similar, your colour choice is ignored and a more visible colour is used. Check this box to disable the colour check.

Axis drag may not invert y axis range

If you check this box, dragging a y axis or the x axis in an XY view to change the scale (a drag starting in the number label region) will not change the axis direction. If you check the box you can still edit the axis limits to invert the sense of the axis by double clicking the axis to open the axis dialog.

Log axis default decades

This field sets the number of decades to display when you swap an axis to logarithmic mode from the X or Y axis dialogs or optimise the Y axis and the low limit is less than or equal to 0. In these cases, the low axis limit is set to be the high axis limit divided by 10 to the power of the number of decades. For example, with 4 decades, if the upper value was set to 400, the lower value would be set to .04 (that is, $400 / \text{Pow}(10, 4)$).

Round dot minimum size

From Spike version 7.01 you can choose to draw round dots in time and result views rather than square dots. This field sets the minimum size (in pixels) at which to do this or you can set the value 0 to always draw dots as squares. Not that printed output will typically have a much smaller pixel size than the screen, so dots that draw as squares on the screen may draw as circles when printed if you enable this option.

Sampling

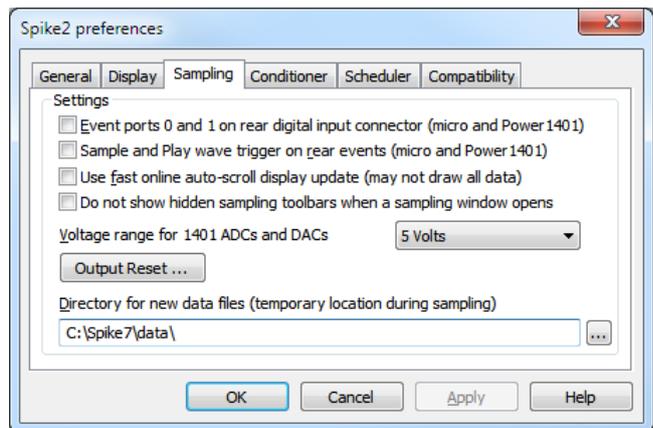
This tab sets preferences associated with sampling data and the 1401.

Event ports 0 and 1 on rear digital input connector

Unless you use a 1401plus, you can source event ports 0 and 1 either from the front panel or from the digital input connector. Check the box to use the rear panel to match the 1401plus. Leave the box unchecked to source your event signals from the front panel Event 0 and 1 inputs.

Sample and Play wave trigger on rear events

Unless you use a 1401plus, you can trigger sampling and on-line waveform output with the front panel Trigger BNC, or on the rear panel Events connector pin 4 (GND is pins 9-15). Check the box to use the rear connector.



Use fast online auto-scroll display update The online automatic scrolling of new data into the display during sampling can make the host computer feel unresponsive, especially in complex display modes at high sampling rates. If you check this box, Spike2 uses a faster (but incorrect) algorithm to decide how much of the screen to repaint each time the view scrolls during sampling.

In the faster mode, Spike2 does not properly allow for changes in already drawn data that depend on newly sampled data that has just scrolled into the display. This is particularly visible with sonograms and channels with channel processing options such as time shifts.

Do not show hidden sampling toolbars when a sampling window opens When you open a data file for sampling, Spike2 will normally also show the Sampling control toolbar, the Sample Status toolbar and the Output Sequence control bar (if there is an output sequence). If you check this box, these toolbars remain in their current state. You may wish to do this to prevent the toolbars changing the screen arrangement.

Voltage range for 1401 ADC and DACs Most 1401s have a ± 5 Volt input range, but some have ± 10 Volts. Spike2 detects this automatically in recent 1401s, but you must set it manually for the 1401*plus*. Choose from 5 Volt, 10 Volt and Last seen hardware. You are warned if Spike2 detects a conflict between user settings and installed hardware. The voltage range affects scaling in the sampling configuration and DAC output values in the output sequencer. It has no effect on scale values in previously sampled data files.

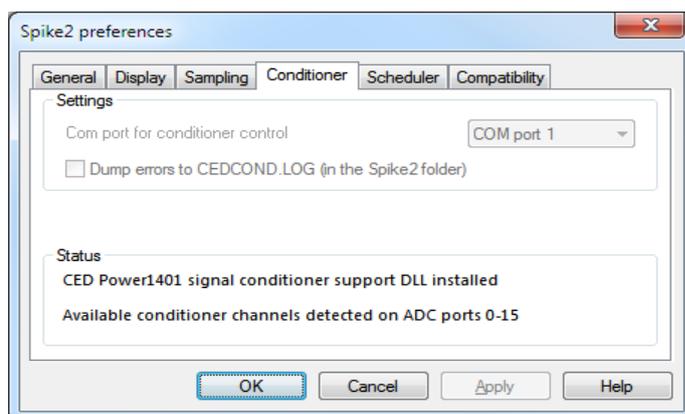
Directory for new data files This is the directory/folder where Spike2 stores files data created by File menu New during sampling. If you do not set a directory, Spike2 uses the current directory, which may not be where you expect, so it is a good idea to set one. If you have more than one disk drive, choose a folder in your fastest drive. Do not use a networked drive. New data files in this folder do not have the .SMR file extension.

When you close a new data file, Spike2 prompts you for the file name. What happens next depends on where you choose to save the file. If the file is on the same volume (disk drive) as the directory/folder set in the preferences, Spike2 just renames the file. If the volume is not the same, Spike2 copies the file, and then deletes the original.

Outputs This opens the Application Output Reset dialog to set the reset state of the DACs and digital outputs of a data acquisition device. The same controls are available in the sampling configuration and normally override these settings. This dialog was added at version 7.11 and it is used to set the output state both before and after data acquisition.

See the Sampling configuration Automation tab for a description of the fields.

Conditioner This tab lets you set the port used by a signal conditioner (see the *Programmable signal conditioners* chapter). Check the *Dump errors...* box to write diagnostic messages to CEDCOND.LOG in the Spike2 folder. If you wish to change the type of conditioner you must run the Spike2 install program; you will be offered a list of conditioners to choose from.

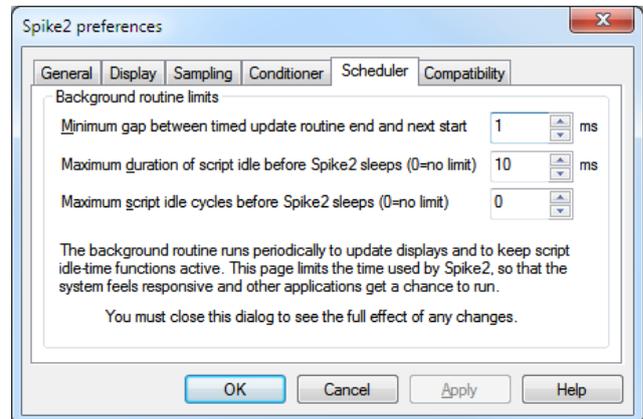


If the currently installed signal conditioner uses a serial line port, you can select any port in the range COM 1 to COM 8. Each time you change the port, Spike2 searches for any conditioner attached to it and will update the status panel.

The Status panel gives information on the type of signal conditioner support that is installed. If it can detect a signal conditioner, it also holds information about the channels that have conditioner support.

Scheduler

This tab limits the processor time consumed by the Spike2 user thread while sampling and idling with a script running. If Spike2 takes too much time, the system feels unresponsive. The tab does not affect the time-critical thread that writes sampled data to disk. You can read more about threads below. Spike2 runs a background routine when it has idle time and also periodically on a timer. The routine handles the following tasks:



- When sampling or rerunning, it gives windows a chance to detect that the maximum time has changed, which may cause windows to scroll and processing to occur. Any invalidated windows will update the next time Spike2 gets idle time.
- If a script is running, it gives any "idle function" in the script a chance to run.
- In automatic file naming mode it starts the next file running.

There are three fields that limit the time used in the background routine. They have no effect on the time used when Spike2 runs a script that does not idle (see the `Yield()` or `YieldSystem()` script commands for this). The standard values work for most cases.

Minimum gap between timed update routine end and next start

If the time interval set by this field passes without the background routine running, it is scheduled to run as soon as possible. You can use values in the range 1 to 200 milliseconds. The standard value is 20 milliseconds. The lower the value, the more time Spike2 will spend on background processing relative to other applications. This field also limits the time that Spike2 will sleep for (see the discussion of threads, below).

Maximum duration of script idle before Spike2 sleeps

When Spike2 gets idle time (see the discussion of threads, below), you can limit the time it uses before Spike2 goes to sleep. Spike2 uses idle time to run script idle routines, such as those created by `ToolbarSet(0, ...)`. You can set from 0 to 200 milliseconds (0 means no time limit). The standard value is 10 milliseconds. The larger the value, the more the script idle routine runs at the expense of other applications.

Maximum script idle cycles before Spike2 sleeps

As an alternative to limiting the script idle routine by elapsed time, you can limit it by the number of times it is called. You can set from 0 to 65535 times (0 means no limit). The standard value is 0. Setting both this and the *Maximum duration...* field to 0 is unkind to other applications. Setting this to 1 is the most generous to other applications.

Threads A *thread* is the basic unit of program execution; a thread performs a list of actions, one at a time, in order. To give you the impression that a system with one processor can run multiple tasks simultaneously, the system scheduler hands out time-slices of around 10 milliseconds to the highest priority thread capable of running. Tasks at the same priority level share time-slices on a round robin basis. Lower priority tasks rely on higher priority tasks "going to sleep" when they have nothing to do or when they are blocked (for example, waiting for a disk read). If this did not happen, low priority tasks would not run.

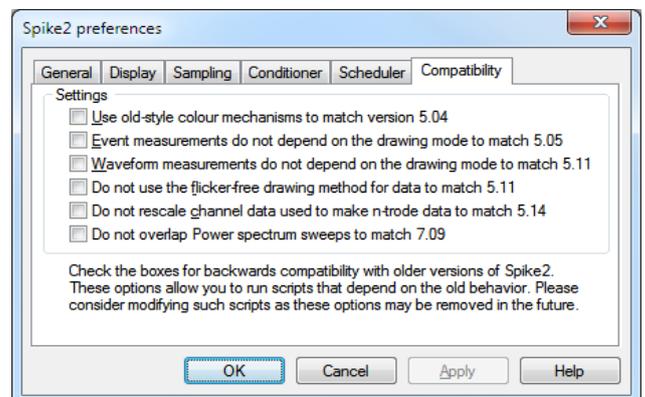
There are also very high priority tasks, usually associated with hardware device drivers, that can interrupt the scheduling system to respond to external events within microseconds. These interrupts usually only last a few microseconds; if a device driver needs a longer period of time it will request a time slice to complete its work and wait for it to be granted.

When Spike2 gets a chance to run, it processes pending messages such as button clicks, keyboard commands, mouse actions and timer events and then updates invalid screen areas. Finally, Spike2 is given idle time until it says it does not need any or new messages occur. If Spike2 needs no more idle time it sleeps until a new message appears in the input queue. The *Minimum gap...* timer wakes up Spike2 if nothing else happens.

Most actions in Spike2 (data display, interactive dialogs and the script system) run in the context of a single user thread, so they have to wait for previous actions to end. Separate threads are used for the real-time sampling thread that moves data sampled by the 1401 into the Spike2 data filing system, for the video capture and replay system and for a thread used to calculate FIR filter coefficients in the filtering dialog. These operations can overlap with the user thread actions.

In a single processor system, the processor swaps around between all available threads, assigning time to them based on their priority. If your system supports multiple processors or multi-core processors or hyper threading, then threads can run simultaneously, and you will feel the benefit of this in smoother and faster system operation, particularly when sampling data or recording video.

Compatibility This tab disables program features that have changed the way that Spike2 works and that might affect your results. We will keep these items for all version 7 releases; we do not promise to keep them beyond this. If you need to check any of the boxes, **please let us know** so we have some idea of the number of users who depend on the old behaviour.



Use old-style colour mechanisms Before version 5.04, the colour palette was saved in sampling configuration files; loading a sampling configuration set the colour preferences. We now save the palette in the registry. Check this box to read colours from configuration files, as before

Event measurements do not depend on the drawing mode Before version 5.06, measurements taken from event channels drawn as mean frequency or instantaneous frequency returned the count of events between two times. Now, the measured values relate to what you see on screen. Check the box for the old behaviour.

- Waveform measurements do not depend on drawing mode* Before version 5.12, measurements from waveform channels returned the nearest data point. Now they return what you see on screen. Check the box for the nearest data point.
- Do not use the flicker-free drawing method* Version 5.12 implements a new buffered drawing method that reduces screen flicker on updates. However, this may impact the display speed. Check the box for the old method.
- Do not check channel scaling and units in n-trode formation* Version 5.15 onwards scales each channel separately when making n-trode data from multiple waveform channels. Check the box to use the first channel scaling for all channels.
- Do not overlap Power spectrum sweeps to match 7.09* From version 7.10, we overlap the Power spectrum sweeps. This will tend to weight the data more equally and can give more consistent results at the cost of more time spent transforming the data. Check this box for the old behaviour.

View menu

This menu controls the appearance of the Spike2 time view, result, XY and text windows. The menu is divided into regions. The first controls the visibility of the main toolbar and status bar. You can also control these by right clicking in the toolbar docking area. The contents of the remaining menu regions depend on the type of the current view.

Enlarge View Reduce View

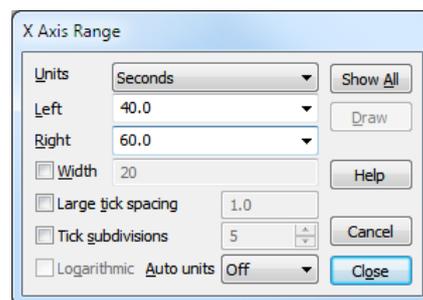


These two commands duplicate the two buttons at the lower left of time and result windows. The enlarge command, short cut `Ctrl+E`, doubles the x axis data region and the reduce command, `Ctrl+R`, halves the region. The left hand window edge is fixed unless enlarging would display data beyond the end of the data, in which case the displayed area is moved backwards. If enlarging would display more data than exists, all the data is displayed. Short cut keys `Ctrl+U` and `Ctrl+I` zoom about the screen centre. You can also change the view size by clicking and dragging the x axis numbers.

X Axis Range



This menu and toolbar command duplicates the action of double clicking on the x axis in a time, result or XY view. In a time view, Units selects Seconds, milliseconds, hh:mm:ss (hours, minutes, seconds) or Time of day. In result and XY views, if the axis units are s or seconds you also have the choice of milliseconds.



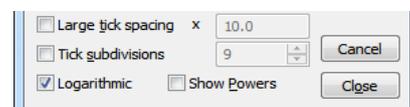
Time of day mode draws times as hh:mm:ss on the x axis that are the time of day at which sampling started plus the time offset into the file. This is for display purposes only; all times used within Spike2 are times in seconds from the start of sampling. All times entered into dialogs are relative to the start of the file. Data files from Spike2 version 4.01 and earlier do not hold a start time and display Time of day the same as hh:mm:ss.

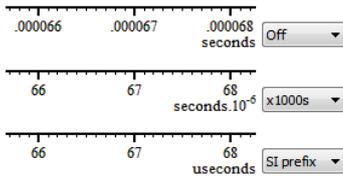
The Left and Right fields set the view start and end in the current units. You can type in new positions or use the drop down lists next to each field. The drop down list contains the initial field value, cursor positions, the minimum and maximum allowed values and the left and right view edges (`XLow()` and `XHigh()`). The Width field sets the window width if the Width box is checked. Click the Draw button to apply changes in these fields to the window. Show All expands the x axis to display all the data and closes the dialog.

In normal use, you will let Spike2 organise the x axis style. However, when preparing data for publication you may wish to set the spacing between the major tick marks and the number of tick subdivisions. If you prefer a scale bar to an axis, you can select this in the Show/Hide channel dialog.

You can control the Large tick spacing (this also sets the scale bar size) and the number of Tick subdivisions by ticking the boxes. Your settings are ignored if they would produce an illegible axis. Changes made to these fields take effect immediately; there is no need to use the Draw button.

The Logarithmic check box changes the x axis to logarithmic mode in result and XY views. In this mode, the large tick spacing field becomes the multiplying factor between large ticks and is usually set to 10. You can use all drawing modes with logarithmic axes; however, straight or cubic-splined lines between points will not pass through the same values as with linear axes. Show Powers is present in logarithmic mode. With this box checked, labels at major ticks display as powers of 10 (or of the value set in the Large tick spacing field).





Auto units is present when Logarithmic is unchecked and Units is not milliseconds hh:mm:ss or Time of Day. Off selects unscaled axis units. SI prefix or x1000s scales the axis units by powers of 1000 where this makes sense to make the axis easier to read. x1000s displays the scale change as $s \cdot 10^{-3}$ for example, SI prefix prepends one of the SI scale prefixes (G, M, k, m, u, n...), for example ms. Scale changes made to the axis are cosmetic only; values used by Spike2 and scripts are in the original units.

Cancel undoes changes made with the dialog and closes it. Close ends the dialog.

Short cut keys

Short cut keys that control the x axis are: Home and End move to the start and end of the data, Left and Right arrow scroll by one pixel, Shift+Left and Shift+Right scroll by several pixels and Ctrl+Left and Ctrl+Right move by half the screen width. The mouse wheel will also scroll the x axis one pixel at a time; add Shift or Ctrl to scroll by larger numbers of pixels.

Jump to event

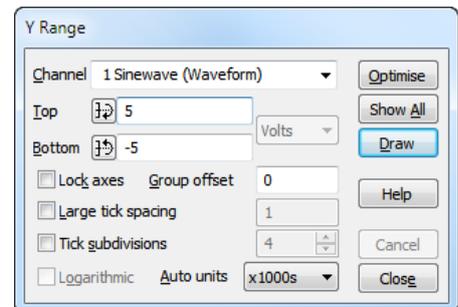
In a time view you can jump to the next or previous event by selecting the event channel, then using Alt+Right/Left arrow. Spike2 searches for the nearest event to the centre of the screen in the specified direction. If more than one event channel is selected, all channels are scanned for the nearest event. Spike2 beeps if no event is found.

You can jump to any TextMark from the TextMark dialog. Double-click any TextMark to open the dialog, click the >> button, select the marker from the list and click Show.

Y Axis Range



This dialog sets the y axis range and style for time, result or XY view channels. The Channel field chooses one, all or selected channels (see the *Getting started* chapter to read about selecting channels). With multiple channels, the Top and Bottom values are for the first channel that matches the displayed channel units.



Click Optimise to draw visible data scaled to use the available vertical space. Show All displays the full data range for waveform channels and from 0 to the estimated event rate for event channels with a frequency axis. Both buttons close the dialog. You can optimise without opening the dialog with the keyboard short-cut Ctrl+Q and by right-clicking on a channel and selecting the Optimise option from the context menu.

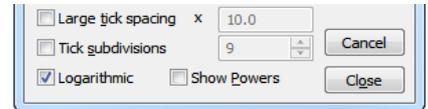
The Top and Bottom fields set the values at the top and bottom of the axis; the units are displayed to the right of these fields. The buttons to the left make the axis symmetric about zero. If there are multiple channels selected, and the channels have different units, you can select the units or <any> to match any units. Click Draw to apply changes to the Top, Bottom, Lock axes and Group offset fields for channels that match the displayed units; select <any> as the units to apply the values to all channels in the list. Cancel undoes any changes and closes the dialog. The Close button closes the dialog.

Lock axes and Group offset are visible when the current channel shares its y axis with other channels and are enabled when the current channel is the first in the group. If you check Lock axes, the grouped channels not only share the same space, they also share the y axis of the first group channel. Group offset sets a per-channel vertical display offset to apply to each locked channel to space out channels with the same mean level.

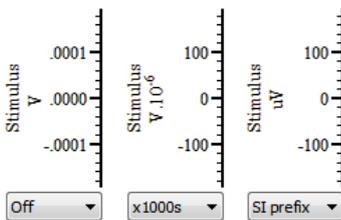
When preparing data for publication you may wish to set the spacing between the major tick marks and the number of tick subdivisions. If you prefer a scale bar to an axis, select this in the Show/Hide channel dialog. You can control the Large tick spacing (this also

sets the scale bar size) and the number of Tick subdivisions by checking the boxes. Your settings are ignored if they would produce an illegible axis. Changes made to these fields take effect immediately; there is no need to use the Draw button.

The Logarithmic check box can be used in Result and XY views to change the y axis to logarithmic mode. In logarithmic mode, the large tick spacing field becomes a multiplying factor between large ticks and Group offset becomes Group factor and must be greater than 0. You can use all drawing modes with logarithmic axes; however, straight or cubic-splined lines between points will not pass through the same values as for linear axes.



The Show Powers check box is present in logarithmic mode and when checked, the axis labels at major ticks display as powers of 10 (or of the value set by Large tick spacing).



Auto units is present when Logarithmic is not checked. Off selects normal axis units. SI prefix or x1000s scales the axis units by powers of 1000 where this makes sense to make the axis easier to read. x1000s displays the scale change as $V \cdot 10^{-3}$, for example. SI prefix prepends one of the SI scale prefixes (G, M, k, m, u, n...), for example mV. Changes made to the axis are cosmetic; values used by Spike2 and scripts are in the original units.

Standard Display

This command sets the current view to a standard state. Time, result or XY view axes are displayed with user options for grids, tick spacing, scaling and special modes turned off. In time and result views, duplicate channels are deleted and all channels display in a standard mode and size in the order set in the Edit menu Preferences, special channel colours are reset, channel processing is removed and the first marker code is displayed. In XY views, all channels are visible, point display mode is set to dot at the standard size, the points are joined and the x and y axis range is set to span the range of the data.

In a text-based view the command removes any maximum line limit except in the Log view, where it applies the line limit set in the Edit menu Preferences command General Tab. It also hides line numbers and displays the gutter and the folding margin (for views that support folding). All the text styles are set to the default values (equivalent to opening the Font dialog and using Reset All) and any text zooming is removed.

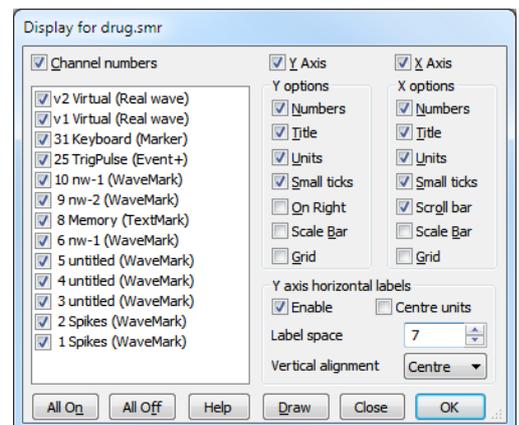
Show/Hide Channel



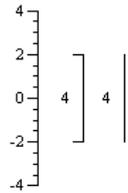
This resizable dialog sets the channels to display in data views. It also controls the display of axes, grids and the horizontal scroll bar in time and result views.

Check the Channel numbers box for channel numbers in time and result windows. The All On and All Off buttons select all or none of the channels, to save time when there is a long list. The Draw button updates the data view. You can show or hide the axes, grid, axis numbers, the big and small ticks and the axis title and units.

You can choose to show the y axis on the right of the data and have horizontal titles and units in place of vertical labels.



For publication purposes, it is sometimes preferable to display axes as a scale bar. If you check the **Scale only** box, a scale bar replaces the selected axis and any grid set for that axis will vanish. You can remove the end caps from the scale bar (leaving a line) by clearing the **Small ticks** check box. The size of the tick bar can set by the **Large tick spacing** option in the **Y Axis Range** or **X Axis Range** dialogs, or you can let Spike2 choose a suitable size for you.



Short cut to display a single channel

In a time or result view, you can double-click a channel with a y axis to hide all other channels; the channel expands to fill the entire window. A second double-click restores the display. **Ctrl+double-click** displays the channel plus all duplicates.

Info

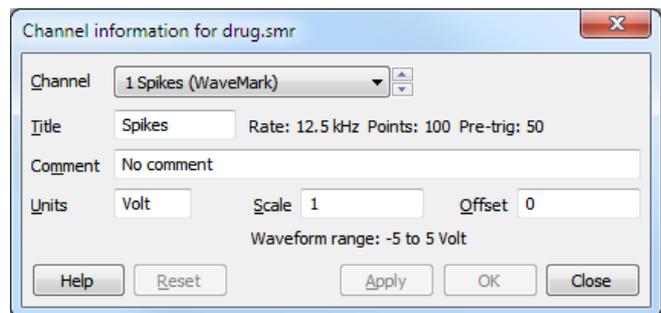
This command displays information about the current result view window. In particular, it displays the number of sweeps that have been added into a PSTH, correlation or waveform average, the number of data blocks in a power spectrum, the number of cycles in a phase histogram and the number of intervals that have been processed to build an interval histogram (including intervals that fell outside the histogram).

File Information

This displays information about the current time view, including five lines of comments and the time and date it was created (if available). You can edit the comments unless the file is open in read only mode. You can cause this window to open automatically when sampling ends from the Automation tab of the Sampling Configuration dialog.

Channel Information

Use this dialog to view and edit time and result view channel information. You can open it with a double-click on a channel title, from the View menu or right click the channel to open the context menu. You can edit the **Title** and **Comment** of the channel set by the **Channel** field.



The remaining fields are hidden or displayed depending on the channel type. The **Reset**, **Apply** and **OK** buttons are disabled until you make a change to one of the fields. The **Close** button closes the dialog and does not apply any changes.

Changes made in this dialog have no effect until you click **Apply** or **OK** (which is the same as **Apply** then **Close**). Changes to all fields except **Title** apply to all duplicates of the channel and to the channel and its duplicates in any window duplicated from the current window. If you change channel without applying changes, any changes are lost.

Changes to the **Title** field are applied to duplicate channels if the duplicate has not been given its own title. If you set the title of a duplicate channel, it has no effect on any other channel. If you clear the title of a duplicate channel, it takes the title of the channel from which it was duplicated.

The **Reset** button restores any changes you have made to the channel settings unless you have used **Apply**. To undo applied changes, close the dialog and use the Edit menu **Undo** command (**Ctrl+Z**).

Scale, Offset and Units In a Time view, for waveform, RealWave or WaveMark channels, extra fields show the *Scale* and *Offset* values and the channel user *Units*. The scale and offset convert between the 16-bit integers used to store waveform and WaveMark data and user units. They also convert RealWave values to integers, when required.

$$\text{Real value in user units} = 16\text{-bit value} * \text{Scale} / 6553.6 + \text{Offset}$$

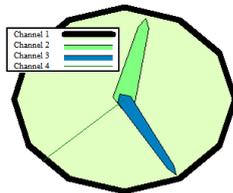
$$\text{Integer value} = (\text{Real value in user units} - \text{Offset}) * 6553.6 / \text{Scale}$$

When the **Scale** and **Offset** fields are present, the **Waveform range** field displays the range of values that a 16-bit waveform channel could span. You can read more about scaling in the documentation for the sampling configuration dialog. Both the scale and the offset must be smaller than 10 billion. The scale may not be set to 0 or very close to 0. If the channel scale or offset is edited into an illegal state, a warning message appears in the dialog and you cannot **Apply** the changes. If you want to calibrate a channel, it is often easier to use the **Calibrate** option of the **Analysis** menu.

In a **Result** view, the *Units* field is always present and can be edited unless the channel is a duplicate.

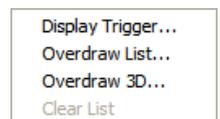
Rescale In a Time view, the **Rescale** button appears for RealWave channels. Click it to set the scale and offset fields so that the full range of data could be represented by 16-bit data. The offset is set to 0 if this does not lose too much precision. Some routines in Spike2 treat RealWave data as 16-bit integer, and the scale and offset you set determine how the conversion from 32-bit real to integer is done.

Options This command is for XY windows only and opens the XY options dialog. This dialog controls the XY window “key”. The key is a small region to identify the data channels that you can drag around within the XY window. For each visible channel it display the channel name and draws the line and point style for the channel. This dialog also has a checkbox that controls the automatic expansion of the axes when new data is added. The equivalent script language command is `XYKey()`.



This example (see `clock.s2s` in the `Scripts` folder) shows the key. You can the key background transparent or opaque and choose to draw a border around the key. If you move the mouse pointer over the key, the pointer changes to show that you can drag the key around the picture. Double-click the key to open the Options dialog.

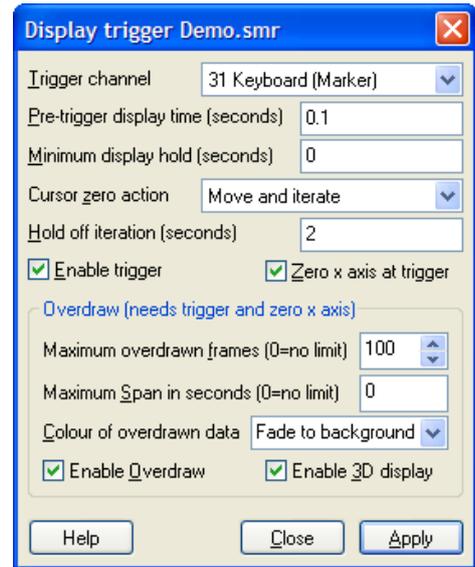
Trigger/Overdraw This View menu command is for time views only and opens a pop-up menu that leads to commands for controlling triggered displays, display overdrawing and 3D (three dimensional) displays of sweeps of data identified by trigger events:



- | | |
|------------------------|---|
| Display trigger | This is the overall setup dialog for triggered and overdrawn displays. Overdrawn displays use a stored list of trigger times. |
| Overdraw List | A quick way to overdraw using an event channel as a source of multiple trigger times. |
| Overdraw 3D | Controls how frames are shifted and scaled to give a 3D effect. |
| Clear List | Empty the stored list of trigger times. |

Display Trigger

The display trigger is used with Time views to provide an oscilloscope style trigger, paged display on-line, display overdraw, a 3D display and a means of easily moving back and forward to the next or previous trigger both on-line and off-line. The script language equivalent of this command is `ViewTrigger()`.



Trigger channel

This field selects an event, Marker, WaveMark, TextMark or RealMark channel as the trigger. You can also select Paged display for on-line paged displays.

Pre-trigger display time

The Pre-trigger display time field sets the time before the trigger to show each time a trigger occurs. This dialog does not set the width of the time view; that is set by the normal time view mechanisms. If you set the pre-trigger display time larger than the displayed time view width or negative, the trigger point will not be visible. Negative pre-trigger times move the trigger point off the screen to the left of the display.

Minimum display hold

The Minimum display hold field is used on-line and sets the minimum time that data is displayed after the current time passes the right-hand edge of the screen. A value of 0 means wait until the current screen is displayed before looking for new triggers.

Cursor zero action

The Cursor zero action field has three setting that control what happens to cursor 0 and any active cursors that depend on it when the view triggers:

- No action Cursor 0 state is unchanged.
- Move to trigger Cursor 0 moves to mark the trigger point, active cursors do not move.
- Move and iterate Cursor 0 moves to mark the trigger point, active cursors move.

Hold off iteration

This field appear when Cursor zero action is set to Move and iterate. It is used online and sets a time, in seconds, to delay the active cursor iteration after the trigger. This allows you to search for data features after the trigger point.

Zero x axis at trigger

The Zero x axis at trigger checkbox changes the x axis so that 0 lies at the trigger time. This is purely a visual convenience; all measurements are still in the original x axis units.

Enable trigger



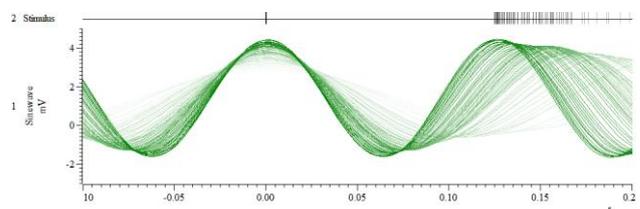
Check the Enable trigger checkbox and click OK to allow display triggering. Two extra buttons appear in the area to the left of the x axis scroll bar. These buttons can be used on-line and off-line to move backwards and forwards by a trigger. With Paged display set, they move by the time between the trigger point and the right hand screen edge.

Keyboard control

In addition to clicking the buttons, you can also step to the next and previous trigger point with the `Alt+Shift+Right` and `Alt+Shift+Left` key combinations.

Overdraw

The overdraw section of the dialog is enabled when both Enable trigger and the Zero x axis at trigger are checked. In overdraw mode, data sections identified by triggers are drawn over each other in time order, oldest first up to the current time. You can change the displayed time range as normal, but



trigger times are ignored if time before 0 or time after the end of the file would be displayed. Each new trigger time is added to a list of previous times. If you step backwards, all trigger times after the time you step to are forgotten. Accepting the dialog settings with **OK** clears the list if the new settings are incompatible with the old settings (for example if the channel is changed).

Overdraw is just a display mode. All measurements, cursor positions and the like behave as if the overdrawn frames are not present.

Maximum overdraw frames You can limit the number of overdrawn frames of data. Values up to 4000 are allowed, or 0 meaning no limit on the number of frames. It takes time to draw each frame, and you will experience significant screen update delays if you display huge numbers of overdrawn frames of data over long time periods.

Maximum span in seconds You can also limit the time range of the overdrawn frames with this field. Each time you add a trigger time, all trigger times after the new time and any trigger time that is more than the Maximum span before are discarded. You can set the value 0 for no time limit.

Colour of overdrawn data The last frame drawn (which corresponds with the trigger time, that is the time for which the x axis is showing 0) is always in the standard colour. You can choose how the remaining frames are drawn from:

No change	Draw in the normal colours.
Half intensity	A equal mix of the normal colour and the background colour.
Fade to background	A gradual fade from the normal to the background colour. In 3D mode, the colour depends on the z axis value, otherwise it depends on the frame number.
Fade to secondary	Same as Fade to background unless a secondary colour is defined, when the fade is to the secondary colour.

Enable Overdraw Check this box to enable overdraw mode. You must also have checked **Enable trigger** and **Zero x axis at trigger**. Frames are added to the overdraw list by stepping to the next trigger event or with the **Overdraw List** dialog.

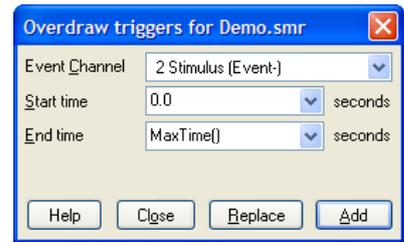
Enable 3D display Instead of overdrawing the frames of data exactly on top of each other, you can choose to draw them offset both vertically and horizontally to create a three-dimensional effect. Check this box to enable this drawing method. The 3D drawing only occurs if you have also checked **Enable trigger**, **Zero x axis at trigger** and **Enable Overdraw**. The screen arrangement of channels and frames in 3D drawing mode is controlled by the **Overdraw 3D** dialog.

On-line and in Rerun When enabled, incoming trigger channel data is searched until a trigger is found when the display will hold with the pre-trigger time shown before the trigger until another trigger occurs. The hold time will be at least as long as the **Minimum display hold** field from the point where the display reaches the right-hand edge of the screen. In **Overdraw** mode, the display hold time limits the number of frames that will be overdraw as if you trigger at time t , the earliest time the next trigger will be used is: $t + w - p + h$ where w is the screen width, p is the pre-trigger time and h is the display hold time.

With **Paged display** selected, each time the right-hand screen edge is reached and the hold time has passed, a new sweep starts. The pre-trigger time sets the overlap between the sweeps.

Overdraw List

In a time view you use this View menu command to add a list of trigger times to a time view based on the settings in the Display Trigger dialog. If the time view is not in overdraw mode, the dialog displays a warning message and a button to open the Display Trigger dialog where you can set overdraw mode. The script language equivalent is `ViewOverdraw()`. Dialog controls are:



Event Channel An Event, Marker or Marker-derived channel in the time view to use as a source of trigger times to add to the list held by the time view.

Start time, End time The time range to search for events to add to the list. All events found in the time range are used as trigger times.

Replace Click this button to clear any existing times from the list before adding new times.

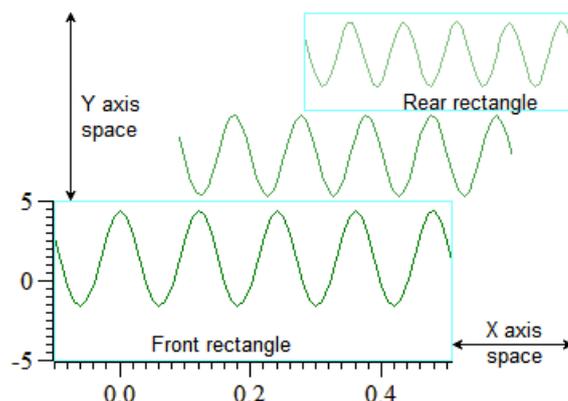
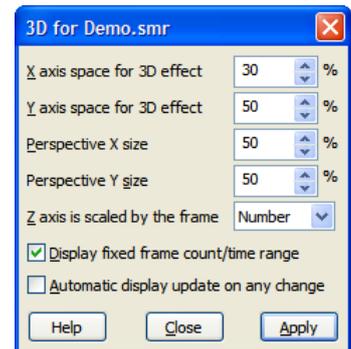
Add Click this to add times identified by the channel and the time range to the list. Unlike manually adding times by stepping to the next or previous event when the last added time causes all later times to be deleted, added times are merged into the list in ascending time order. The current time display is set to be the last time in the list.

Display Trigger dialog interaction

If you add events to the trigger list such that the event times come from a mix of channels (or if you add times not associated with a channel using a script), the Display Trigger dialog channel will change to No channel is set. If you replace events such that the trigger list holds times from a single channel, the Display Trigger dialog will show that as the current channel.

Overdraw 3D

In a time view you use this View menu command to open the 3D dialog to control the 3D drawing effect. The script language equivalent is `ViewOverdraw3D()`. This dialog does not enable 3D drawing; use the View Trigger dialog to do that. To create the 3D effect, each frame is drawn inside a rectangle and the position and size of the rectangle is changed based on a notional z axis. The z axis is based on the frame number, or based on the trigger time of the frame. The z axis can be set to a constant frame count or time range, or it can vary depending on the current set of frames to display.



In this example, there are three frames of data. The front rectangle, used for the last trigger, is always positioned at the bottom left of the channel area. The rear rectangle, used for the oldest trigger time, is always positioned at the top right of the channel area. The middle frame also has a rectangle that is calculated by a linear interpolation between the front and rear rectangles based on the z axis position of the frame.

The dialog fields control the positioning of the front and rear rectangles:

X axis space for 3D effect This field sets the percentage of the width of the channel area to use for the 3D effect. The larger the value, the smaller the width of the front rectangle. Set this and the *Perspective X size* fields to 0 to make all frames draw vertically aligned.

Y axis space for 3D effect This field sets the percentage of the entire view vertical space to share between all the channels to generate the 3D effect. All channels are given exactly the same space so that the 3D effect is the same for all channels. The larger the value, the smaller the height of the front rectangle.

Perspective X size This sets the width of the rear rectangle as a percentage of the width of the front rectangle in the range 0 to 100. Setting a value less than 100 gives a perspective effect.

Perspective Y size This sets the height of the rear rectangle as a percentage of the height of the front rectangle in the range 0 to 100. Setting a value less than 100 gives a perspective effect.

Z axis is scaled by the frame... The position of each frame of data can be set either by the frame number (giving equally spaced frames), or by the frame trigger times. Choose from **Number** or **Time**.

Display fixed frame count/time range If you check this field and you have a set a maximum number of overdrawn frames (for z axis scaled by **Number**) or a maximum time span (for z axis scaled by **Time**) in the **Display Trigger** dialog, the z axis will be of a fixed size. If you do not check this box, or there is no maximum set for the number of frames or the time range, the length of the z axis is set by the number of frames in the list or by the time span of frames in the list.

Automatic display update on any change If you check this box, and change made to the dialog will cause the display to update. If your display does not take long to update you probably want to check the box, but if you have a lot of data to draw you may prefer to update with the **Apply** button.

Notes The 3D display mode behaves in exactly the same way as the **Overdraw** mode except that the overdrawn frames are offset and scaled. There is one other difference; if there are no frame trigger times in the list, then nothing is displayed. There is no restriction on the drawing modes you can use (other than your own common sense). **Overdraw WaveMark** and **Sonogram** modes are unlikely to be useful.

The 3D display mode makes no difference to any measurements you may make. There is still a current time range that is displayed (this corresponds to the frame of data that is at the front of the display).

Clear List This command from the **View menu Overdraw/Trigger** pop-up menu removes all stored times from the list used for overdrawing.

Channel Draw Mode



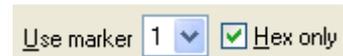
This menu item is available when the current window holds a Spike2 data document or result view and sets the data channels display mode. You can set the mode for a single channel, all channels, or selected channels. The dialog changes, depending on the channel type and display mode. The Draw button is common to all modes and updates the display without closing the dialog box.

Time view drawing modes

In some modes the Edges field appears for event level data. You can select Rising, Falling or Both edges of the data. If you select both edges, then the display modes that show frequency count both the rising and falling edges of the event signal in their rate calculations. You would normally count only one edge, so select the edge you prefer.



For Marker, RealMark, TextMark and WaveMark channels, the Use marker and Hex only fields may appear. The Use marker field appears if the code is displayed or sets the colour of the items. Normally code 1 is used, but you can choose any of the four codes. If you select a code other than 1, the channel number field will display in red as a warning of a non-standard display mode.

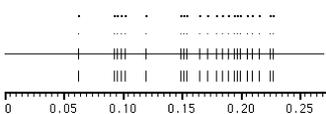


The Hex only field appears if the marker code is displayed. Check the box to display all codes as two hexadecimal digits, unchecked to display codes 0x20 to 0x7e as single characters.

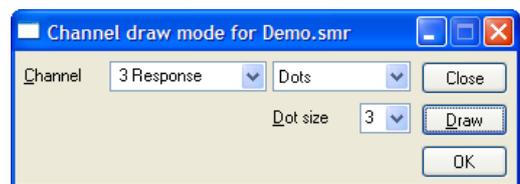
If the display will show the data as dots, the Dot size field appears and you can set any dot size from 0 (the smallest mark possible) to 1 to 9 times the thickness of the lines set for drawing data in the Display tab of the Edit menu Preferences. If you set a size to match or exceed the Round dot minimum size field in Edit menu Preferences, the dots will be drawn as circles. Round dots take longer to draw than square dots.



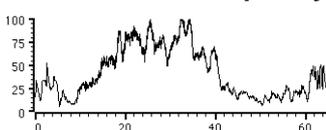
Dots and Lines mode



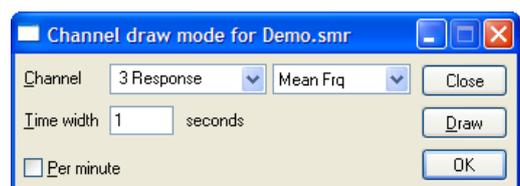
The simplest event channel draw method is dots. You can also select Lines in place of Dots. The picture to the left shows the result of both types of display. If you select lines for a marker channel, the display of marker codes is suppressed. You can also select Dots mode for a waveform channel. In Lines mode you have the Centre line option to show a horizontal line.



Mean frequency



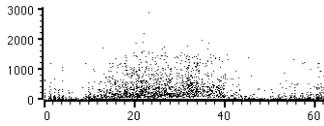
The mean frequency is calculated at each event by counting the number of events in the previous period set by the Time width field in seconds. The result is measured in units of events per second unless the Per minute box is checked. The mean frequency at the current event time is given by:



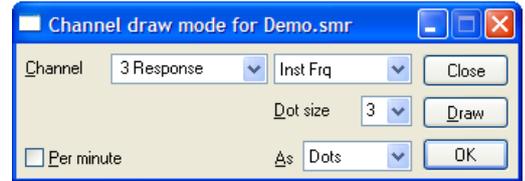
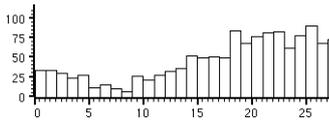
$$\frac{(n-1)}{t_e-t_1} \quad \text{if } (t_e-t_1) > t_b/2$$

$$\frac{n}{t_b} \quad \text{if } (t_e-t_1) \leq t_b/2$$

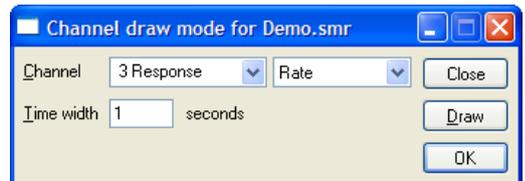
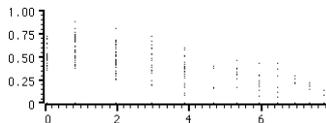
where: t_b is the bin size set t_e is the time of the current event
 t_1 time of the first event in the time range n is the events in the time range
 A constant input event rate produces a constant output until there are less than two events per time period. You should set a time period that would normally hold several events.

Instantaneous frequency

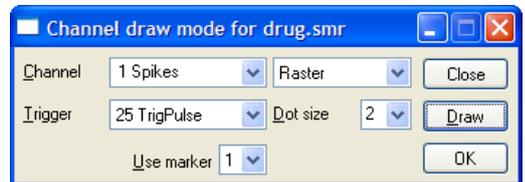
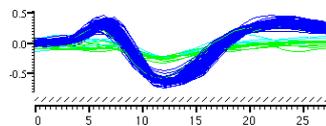
Instantaneous frequency is the inverse of the time interval between an event and the previous one on the same channel. Check the **Per minute** box for a rate per minute rather than per second. You can display the result as **Dots**, **Line** (linking the dots) or **Skyline** (horizontal lines between dots). In **Dots** mode you can choose the dot size. You can display the frequency as events per minute rather than per second.

**Rate histograms**

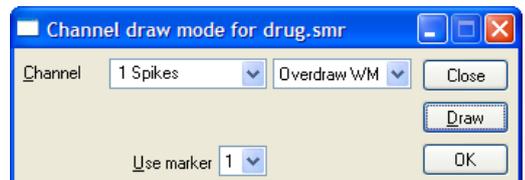
Rate mode counts how many events fall in each time period set by the **Time width** field, and displays the result as a histogram. The result is not divided by the bin width. This form of display is especially useful when the event rate before an operation is to be compared with the event rate afterwards.

**Raster display**

Raster mode shows the event positions relative to trigger times. Each trigger event defines time 0 in the y direction for the sweep. The **Y Range** dialog sets the time range to display in the y direction; negative times show pre-trigger events. For each trigger, events are drawn no further back in time than the previous trigger and no further forward in time than the next. If the channel is marker-based, events are drawn in the colours set for the **WaveMark** codes.

**Overdraw WM**

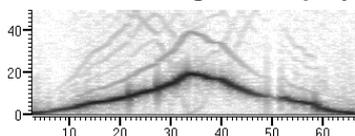
This mode draws **WaveMark** data as superimposed waveforms. Channels in this mode are moved to the top of the window and separated from the x axis (which does not apply to them) by a grey bar. If this mode is used during data capture and the screen is scrolling to show the latest data, new **WaveMark** events are added, but old events are not erased (to reduce flicker). Click on the x axis scrollbar thumb to force an update.

**Find with cursor 0**

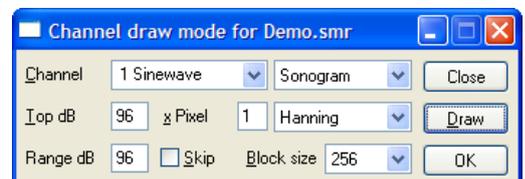
To locate an event, right-click on its waveform at a point where it is clear of all other waveforms and select **Find with cursor 0** in the pop-up menu. This locates the event nearest to the clicked position and moves cursor 0 to it. If the **Edit WaveMark** dialog is open, this will display the event. Any active cursors will iterate.

Set WaveMark codes

Hold down **Alt+Ctrl** and click and drag a line over the events you want to identify. On mouse up, a dialog opens in which you can set codes for the intersected events.

Sonogram display

Sonogram mode shows how the frequency content of a waveform channel changes with time. The y axis units are **Hz** (frequency) and useful results are available for the frequency range 0 to one half of the sampling rate for the waveform channel. By default,



intensity of the frequency content is indicated by a grey scale, the darker the image, the more intense the signal. However, you can choose a colour scale, or create your own scale with the View menu Sonogram Colour command. You can set:

- Top dB** This sets the signal intensity that corresponds with the darkest grey scale. dB means decibels, which is a logarithmic measure of ratio, usually of amplitudes or power. 20 dB is a factor of 10 in amplitude. Spike2 stores waveform data as 16-bit integers, and we measure the amplitude with respect to 1 bit, so 96 dB is the maximum possible level.
- Range dB** This sets the range of data to display as a grey scale. Signals with an intensity of Top dB - Range dB (or less) are displayed as white. If you are unsure what dB values to set for a new signal, setting values of 96 dB for both Top dB and Range dB will display something in almost all cases!
- x Pixel** You can speed up drawing, at the expense of resolution, by setting this field to values greater than 1. It sets the number of screen pixels in the x direction to calculate at a time. A value of 1 gives the best visual resolution (and the slowest calculation and drawing time).
- Window** The sonogram is calculated using a Fast Fourier Transform. As explained in the analysis section for the Power spectrum, it is important to apply a “data window” to the signal before taking the power spectrum, otherwise the results are very difficult to interpret. We provide several different types of window: None, Hamming, Hanning, Kaiser 30dB, Kaiser 40dB, Kaiser 50dB, Kaiser 60dB, Kaiser 70dB, Kaiser 80dB and Kaiser 90dB.
- All windows are a compromise between increasing the apparent width of a spectral peak and the ability to see small signal peaks in the presence of large ones. If you apply no window, you will get the sharpest resolution of a single peak. However, you will not be able to see any small peaks around it due to the “sidelobes” of the window. If you are not familiar with the use of windows, the Hanning window is a reasonable compromise. The Kaiser *xxdB* family of windows has the property that the largest sidelobe is *xxdB* below the peak. Of course, the larger the *xxdB*, the wider the peaks become.
- Block size** This determines the number of data points used in the FFT, and thereby it determines the frequency resolution. Like the choice of data windows, this is also a compromise. The larger the block size, the better the frequency resolution, but the worse the time resolution. If you are looking for a short, localised burst of changing frequency, you will need to use a block size that is smaller than the duration of the episode you are looking for.
- Skip** If you are analysing a lot of data, there can be thousands of data points for each screen pixel. If you check this box, the sonogram will only analyse the first Block size points for each pixel (normally the sonogram analyses all the points). This can save a lot of time if you have a really large file. Of course, the result will only represent a “sampling” of the correct response.

You can export Sonograms as enhanced metafiles and as bitmaps. Sonogram printing is supported with Postscript compatible printers. With other printer types, the intensity scale output may be quite coarse. You may obtain better output by saving the sonogram as a bitmap and printing from a bitmap editing program.

**Waveform, SkyLine,
WaveMark and Cubic
spline**

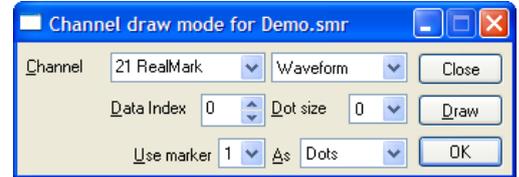
These modes apply to waveform, RealMark and WaveMark channels. Waveform mode joins the data points with straight lines. Skyline joints points with horizontal and vertical lines (not WaveMark channels). Cubic spline mode joins the points with smooth cubic curves based on the assumption that the first and second derivatives of the data are continuous at the data points. However, you must always remember that the only data values that you can rely on are those at the sample points. Cubic spline mode becomes waveform mode in Windows metafile output. WaveMark is for WaveMark channels only and is the same as waveform mode but also draws the selected marker code.

Set WaveMark codes

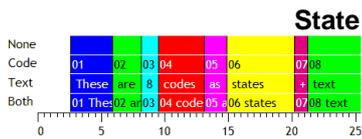
If the channel holds WaveMark data you can identify events that cross a line. Hold down **Alt+Ctrl** and click and drag a line over the events you want to identify. On mouse up, a dialog opens in which you can set codes for the intersected events. If the channel holds multiple traces, all traces are used to select events.

RealMark data in Waveform mode

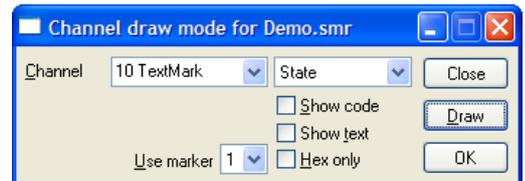
RealMark channels have extra fields in Waveform mode. The Data Index field sets the attached value to display; value indices start at 0. The As field specifies how to connect the waveform points. If Dots is selected, the Dot size field can be used. The other As options are Line, Default (same as Line) and Skyline (which is the same as selecting Skyline as the draw mode).



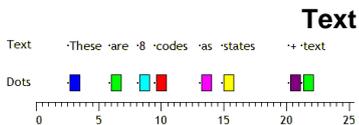
Default (same as Line) and Skyline (which is the same as selecting Skyline as the draw mode).

**State**

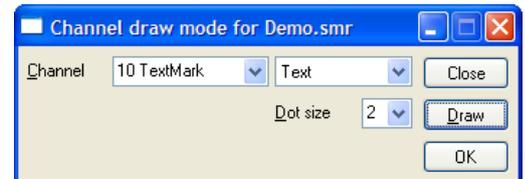
This drawing mode can be applied to a marker, TextMark, RealMark or WaveMark channel. The state is set by the selected marker code and persists up to the next marker on the channel. States are drawn in the same colours as set for the WaveMark codes except that code 00 is drawn in the channel background colour. The initial state at the start of the file and before any markers are found is assumed to be 00. You can choose to display the marker code and if this is a TextMark channel, the text.



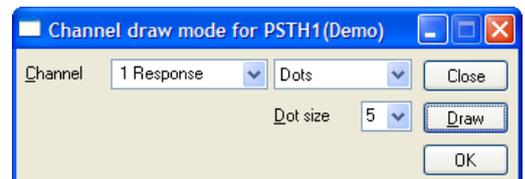
If you set a marker filter on a channel drawn in state mode, this will extend states past the filtered out markers, which may not be desirable. If a channel is in State mode, the Cursor Values dialog reports the displayed state.

**Text**

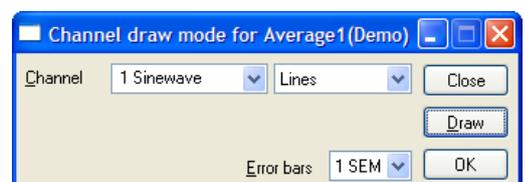
This display mode applies only to TextMark channels. Each item draws as a dot at the TextMark time followed by the text. You can override the dot and text colours with the channel primary and secondary colours in the colour palette. You might also consider using this channel as a Vertical Marker.

**Result view drawing modes**

There are seven drawing modes for result views: Histogram, Line, Dots, SkyLine, Cubic Spline, Raster and Raster lines. The last two can be used if you checked the Raster data box when you created the result view. In Dots mode you set the size of the dots.

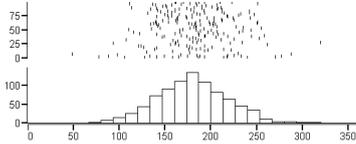
**Error bars**

If your result view has associated error information, for example a waveform average with error bars enabled, you have an extra control. You can choose from None, 1 SEM, 2 SEM or SD. It should be emphasised that error bars only have meaning if the data points that contribute to the average have a normal distribution about the mean. Given this, then 1 SEM shows ± 1 standard error of the mean, 2 SEM is ± 2 standard errors of the mean and SD is ± 1 standard deviation.

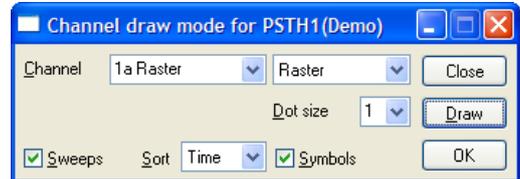


If each point of your data can be modelled as a constant "real" value to which is added normally distributed noise with zero mean, then you would expect the measured mean value to lie within 1 *standard error of the mean* (SEM) 68% of the time, or within 2 SEM 95% of the time. The *standard deviation* represents the width of the normal distribution of the underlying data at each data point.

Result view raster



If a result view channel has associated raster data, there are two more drawing modes: Raster and Raster lines. Raster mode shows each event as a dot, Raster lines mode shows each event as a short vertical line. The Dot size field controls the dot size and vertical line length. In Raster lines mode you can check the Centre line box for a horizontal line through each sweep.



OK

Extra features with auxiliary values

If you set an auxiliary channel in the peri-stimulus, event correlation or phase histograms setup dialog or set auxiliary values from the script language, you can choose the sweep sort order. The Sort field can be set to Time or Sort 1 to Sort 4. Time presents the sweeps in the order that they were recorded. Sort n sorts the sweeps based on the sort value n. Sort value 1 is set if you set an auxiliary channel in the result view setup dialog. The script language can set all the sort values with the `RasterSort()` command.

If you check Sweeps, the y axis is a sweep count and all sweeps are evenly spaced in the y direction, otherwise it is the value of the item selected by the Sort field and the sweeps are spaced out based on the value set in the Sort field. If you sort by Time, the sweep number is the order in which sweeps of data were added into the result view.

Each raster can store 8 times and display them as symbols (circle, cross, square, up triangle, plus, diamond, down triangle, filled square) if the times lie within the x axis range. Symbol time 1 (circle) is set if you select an event channel as the auxiliary channel in the result view setup dialog. Use `RasterSymbol()` from the script language.

Interrupting drawing

Despite our best efforts to draw huge data files quickly, drawing can take a long time. You can interrupt it with the `Ctrl+Break` key combination. The Break key is usually at the top right of the keyboard and is often labelled Pause with Break on the front. If your system has sound enabled and you have selected a sound for "Exclamation", Spike2 plays this sound to confirm that drawing has been abandoned for the current channel.

XY Draw Mode

This sets the drawing style of an XY window. OK applies changes and closes the dialog, Apply makes changes without closing. Cancel closes and ignores any changes since the last Apply. The Channel field sets the channel to edit, or select All channels. If you change channel, the dialog saves any alterations; there is no need to use Apply before changing channel unless you want an immediate update. The other fields are:



Join style

- Not joined
- Joined
- Looped
- Filled
- Fill and frame
- Histogram

There are six join styles. In Not joined style, no lines are drawn between data points. In Joined style, each point is linked to the next by a straight line. In Looped style, the points are joined and the final point is linked back to the first point. In Filled style, the data points are not joined, but they are filled with the XY channel fill colour. In Fill and frame style, the first and last points are joined and linked by a line and the channel is

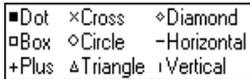
filled. In Histogram style, the points are drawn as a histogram. The Line type and Width fields set the kind of line that joins the points.

Line type and Width



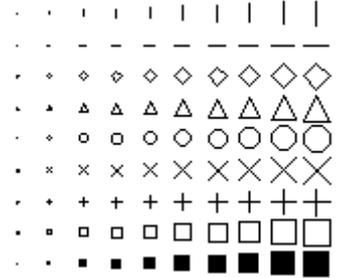
These two fields set the type of line joining data points. The Width field determines how wide the line is in units of half the data line width set by the Edit menu Preferences dialog. Set 2 for the normal line width, 1 for half this width. If the width in pixels is greater than 1, the Line type field is ignored and the line is drawn as a solid line.

Marker and Size



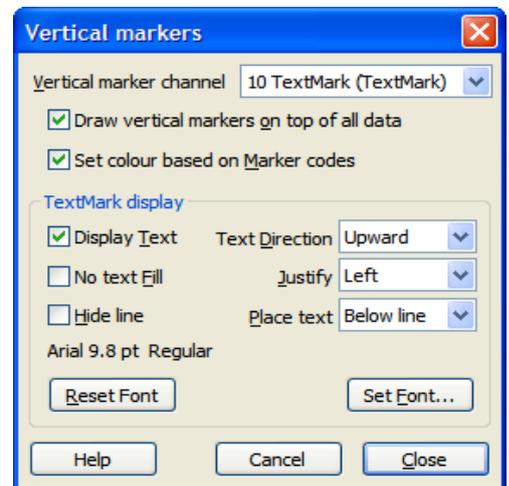
The Size field sets the marker size in units of Points (a point is $1/72^{\text{nd}}$ of an inch or approximately one screen pixel). A size of 0 makes the markers invisible. You can set sizes up to 100 points. There is a wide range of marker styles to choose from.

The picture to the right shows a screen dump of all the marker styles in sizes 1 to 10. If you need to tell them apart on screen, sizes below 3 should be avoided. If you have excellent eyesight and a high-resolution printer, size 1 is viable in printed output.

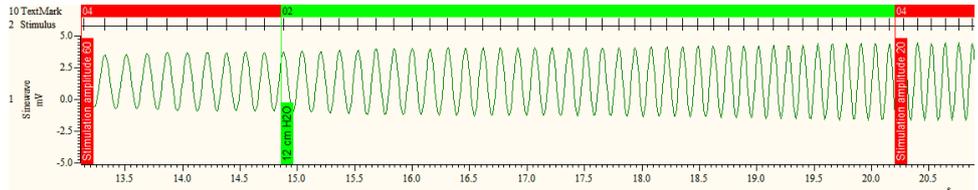


Vertical Marker

The View menu Vertical Marker command is available in a Time view and opens a dialog to control the display of vertical markers. Vertical markers are vertical lines drawn under or on top of the channel displays, rather like vertical cursors, except that their positions are determined by the times of event or marker data points on a channel. If the channel is a TextMark, you have the option of displaying the text of the marker in a variety of styles. Vertical markers are not drawn if the display is in 3D overdraw mode, nor are they drawn over the Overdraw WM area.



TextMark channel draw as State and Vertical marker



Vertical markers are often used in place of vertical cursors where you need a marker that cannot be moved by the user. If you use a memory channel as the vertical marker channel, you can edit the position by adding and removing items in the memory channel.

You can set the width of the vertical line in the Edit Menu Preferences dialog and the default colour in the View menu Set Colours dialog in the Application Colours section.

The dialog fields are:

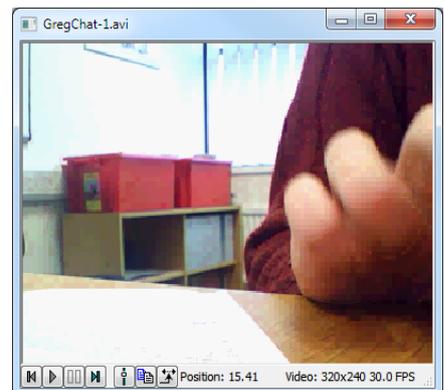
Vertical marker channel

You can select any channel that holds event, Marker, RealMark, WaveMark or TextMark data. You can also select No channel to disable the display of vertical markers.

- Draw vertical markers on top of data* Check this box to draw the vertical markers after the channel data. Leaving this box unchecked places the markers on top of any background (including channel background colours, bitmaps and grids) but under the data. Drawing is faster if this box is checked if you have grids or individual channel background colours defined.
- Set colour based on Marker codes* The colour palette has 9 colours that are used to display WaveMark data, based on the marker code. These colours can also be used for vertical markers if the marker channel is based on a Marker data type. Check the box to enable the use of WaveMark colours.
- TextMark display* The remaining dialog fields are used for TextMark data. If the channel is not a TextMark, these fields are disabled unless **No channel** is selected, when all fields can be set.
- Display Text* Check this box to display the text stored in a TextMark data channel.
- No text Fill* Normally, the area behind the text is drawn in the colour set for the vertical marker and the text is drawn in white or black, whichever provides the most contrast. Check this box to draw the text in the vertical marker colour with no background fill.
- Hide line* When text is drawn, you can choose to hide the vertical line by checking this box.
- Text Direction* You can choose to run the text upwards (to match the text for channel labels on channels with a y axis) or downwards.
- Justify* The text can be **Left**, **Centre** or **Right** justified within the scrollable area of the view.
- Place text* The text can be aligned to be **Below**, **On** or **Above** the vertical marker line.
- Set Font...* Click this button to open a Font dialog where you can choose the font for the vertical marker text. The current font, size and style displays above the **Reset Font** button.
- Reset Font* Click this button to revert the vertical marker font to the font set for the view. The text above this button will show **Use font set for view** if the font has been reset.
- Cancel* The view changes dynamically to show the effects of any changes. If you close the dialog with the **Cancel** button, all changes made are removed before the dialog closes.
- Close* This closes the dialog, accepting any changes. If you have made a change and recording is enabled, script commands are saved that reflect the new state. You can use **Undo** to remove changes after the dialog closes.

Multimedia files

This command opens any multimedia files that are associated with the current time view. It is disabled unless there are suitable multimedia files, identified by their names and location: `fname.smr` has multimedia files in the same folder with names: `fname-N.avi`, where the `N` stands for 1 to 4. If the command is enabled, but no windows open when you use it, the data in the files is corrupt, uses a video or audio codec that is not installed or your computer is not equipped to replay the file.



If a multimedia file only holds audio, the video display area is hidden. The multimedia position is displayed at the bottom of the window and is either the right-hand edge of the time view, or the cursor 0 position. If you use the View menu **Rerun** command, the multimedia windows will keep track of the rerunning data and any audio data will replay through your sound card.

The control bar at the bottom of each multimedia window holds information about the content of the file and control buttons. The displayed frame rate information is the frame rate set in the file. The actual frame rate may be less due to dropped frames because the `s2video` application was commanded to reduce the frame rate or because the recording system could not keep up. The buttons are:



These buttons from left to right step one frame left, replay the file from the current position to the end, stop playing, step one frame to the right. The frame stepping buttons are hidden if the file does not contain video.



This button links the displayed multimedia position to cursor 0 in the associated time view. Cursor 0 is made visible, if it is hidden. With this button down, if you drag cursor 0, the media file moves with it. If the button is up, the multimedia position is the right-hand edge of the time view. This button and the run button are mutually exclusive.



This button copies the video image at the current multimedia file position to the clipboard. This button is hidden if there is no video in the file. The image is copied at the native resolution of the multimedia file, not at the current screen size.



This button sets a time offset between the multimedia time and the associated time view time. This time is saved in the AVI file and can also be set by the `s2video` program.

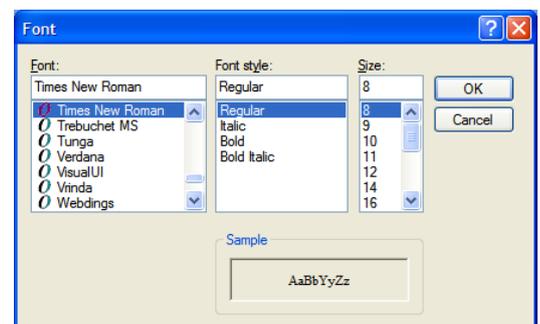
Script users have additional controls over the multimedia window. For example, they can replay individual multimedia windows independently of the associated time view. They can also grab the video image in a variety of formats (RGB, monochrome, HLS) for further script processing. See the `MMOpen()` script command for more information.

A multimedia file contains audio and/or video information. In most cases, the files will have been recorded by the `s2video` application. However, there is nothing to stop you renaming any `avi` file to match your data document. Spike2 requires DirectX to be installed on your computer for this feature to work. For correct operation we strongly recommend that you get the latest release of DirectX from the Microsoft web site. As I write this in September 2003, the latest release is DirectX 9.0b.

Multi-media use makes heavy demands on the hardware of your computer. The more powerful your hardware, the better the result you will get.



You can select the font that is used for each window in Spike2. The font size changes the space allocated to data channels in a time, result or XY view. Smaller fonts give more space to the channels, however fonts need to be large enough to read! You can set different fonts in each data or text window. In a text-based view, this command opens the editor settings dialog for the current text view type where you can set fonts for all the text styles supported by the view. This is described in the Edit menu Preferences under the General tab.

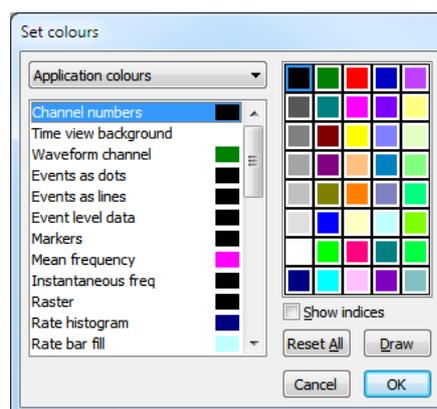


There are other places in Spike2 where you can set fonts for specific items, for instance in the Vertical markers dialog and in the Printed Header and Footer dialog. These open a very similar dialog.

Change Colours



You can choose the colours that are used for almost everything in Spike2. The dialog has separate pages for Application and Marker colours plus additional pages for View, Channel primary, secondary and background colours if a time, result or XY view is current. You select pages with the drop-list at the top left. To change colours, double-click an item in the list on the left to open a colour picker or select one or more items in the list and click a colour in the palette on the right.



Application colours

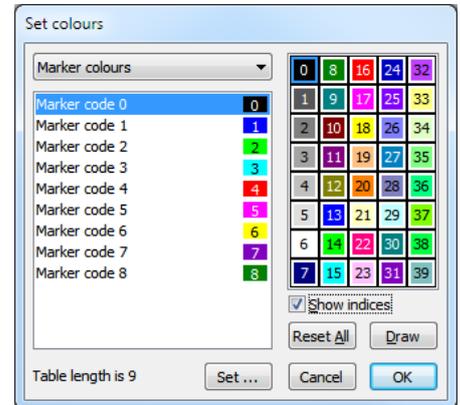
These colours apply to all views. Use Draw to check the result of your action. OK accepts changes, Cancel undoes them; both close the dialog. Reset All returns the list and palette to standard colours. Application items are:

Channel numbers	The channel number. Used to identify, select and drag channels.
Time view background	Background colour of the entire time view (view background)
Waveform channel	Waveform data in time views (primary). Also used for waveforms in the spike shape and digital filter dialogs.
Events as dots	Event data drawn as dots (primary).
Events as lines	Event data as vertical lines (primary), horizontal (secondary).
Event level data	Level event data (primary).
Markers	Marker codes and associated dot, RealMark waveform (primary).
Mean frequency	Events draw as mean frequency (primary).
Instantaneous freq	Events drawn as instantaneous frequency (primary).
Raster	Time view and Result view raster responses (primary), Time view trigger event and Result raster centre line (secondary).
Rate histogram	Border colour for events drawn as a rate histogram (primary).
Rate bar fill	Fill colour for events drawn as a rate histogram (secondary).
Result background	Result view background (view background).
Result lines	Result view data as lines (primary).
Result dots	Result view data as dots (primary).
Result skyline	Result view data as a skyline(primary).
Result histogram	Border of result view data as a histogram (primary).
Result bar fill	Fill of result view data as histogram (secondary).
TextMark text	TextMark data drawn in Text mode
Cursors	Horizontal and vertical cursors and cursor labels.
Controls	This is currently unused.
Grid	All axis grids.
Axes	All axis lines, ticks and labels.
XY view background	The XY view background colour (view background).
Not saving to disk	For new data drawn at the right edge of a sampling window that is not being written to disk. Also used for non-triggered data in the spike shape dialog. Set the same as Time view background for the normal colours.
Errors (SD/SEM)	All result view error bars and envelopes (secondary).
Curve fit	Fitted curves.
Cluster background	The background colour for the clustering dialog.
WaveMark background	The background colour for all the spike shape dialogs.
Vertical markers	Vertical marker colours if marker colours are not used.
XYChannel data	XY view data channels (primary).
XY Channel fill	XY view fill colour (secondary).
XY key	XY view key

If you set a channel or text colour that is too similar to the background, Spike2 will change the colour to keep the item visible. There is an option to defeat this in the Edit menu Display Preferences. The script equivalent is the `ColourSet()` command.

Marker colours Several drawing modes allow you to draw Markers and marker-derived data (TextMark, RealMark, WaveMark channels) in different colours, depending on a marker code. A marker code can take values from 0 to 255.

Marker code 00 is always assigned marker colour 0. All other marker codes are assigned colour $(c \bmod (m-1)) + 1$, where c is the marker code, m is the total number of markers and \bmod is the modulus (remainder) operator. This means that all marker codes apart from 00 are assigned a repeating cycle of codes with marker code 1, $m+1$, $2m+1$, and so on sharing the same colour.



Set Click the **Set...** button to change the number of marker colours in the range 3 to 255. If you increase the number of marker colours, the new colours will start set to black. Script users can set the marker colours with the `ColourSet()` command.

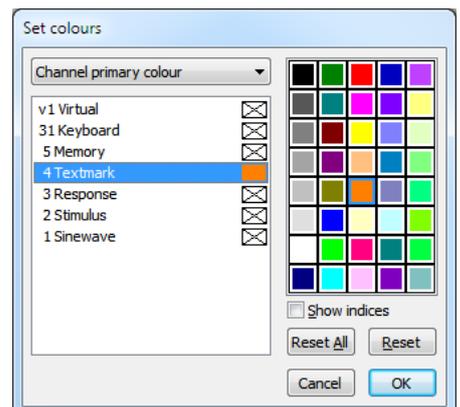
Show indices Check this box to display the palette number for each colour in the palette and the index of items in the list (where appropriate). This is a convenience for script programmers.

Channel primary colour These pages assign colours to data channels in the current time, result or XY view to override Application colours. The primary and secondary colours override items marked (primary) and (secondary) in the Application colour table, above. The channel background colour overrides the view background for the area occupied by the channel data. An X in a box marks a channel with no override.

Channel secondary colour

Channel background colour

Changes made on this page are applied immediately, so there is no **Draw** button. You can **Reset** the selected channels back to the standard colours set in the Application colours page. Channel colour overrides are stored in the resource file associated with the data file and in the sampling configuration file for new data files. Script users can set channel colours with the `ChanColourSet()` command.



View colours This page allows you to override the application colours set for the current view. At the moment you can only override the view background colour. The equivalent script command for this is `ViewColourSet()`.

Changing the Palette



To change a palette colour, double click it to open the colour picker and select a replacement colour. The first seven palette colours form a grey scale from black to white and cannot be changed. You can replace the palette colour with any standard colour, or you can click the **Define Custom Colours...** button to select an arbitrary colour. Click **OK** or **Cancel** to exit.

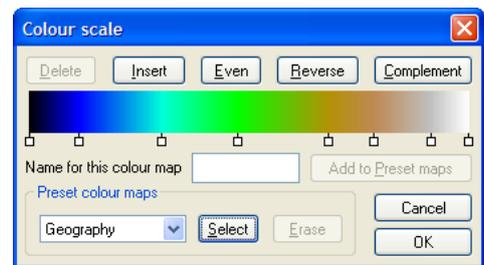
The colour palette is owned by the Spike2 application, not by the data file and is stored in the registry and also with the Spike2 configuration in `default.s2c`. When Spike2 starts, the palette, Application and Marker colours are read from the registry; if there are no colours to read, the colours in the configuration file are used.

Use Colour and Use Black And White

If you have a monochrome monitor or if you have to use a printer which does not handle colour well, you can choose to display your data files in black and white. The **Use Black And White** menu item switches from colour to black and white displays. If you change to monochrome, the menu item changes to **Use Colour**. This option saves you from the tedious task of changing every colour in the colour palette manually.

Sonogram colours

The Colour Scale dialog can be opened by the Sonogram Colours command in the View menu and from the clustering dialog View menu Density Colour Map command. You can also open this dialog by right-clicking on a time view sonogram display and selecting Sonogram Colours from the context menu. The dialog sets the colour gradient for the time view sonogram display mode. You can choose from a range of built-in and user-defined colour maps and create new maps. Apart from the standard **OK** and **Cancel** buttons, the dialog is in three sections: an upper part where you create new maps, the centre line, where you name and save created maps, and the **Preset colour maps** section where you select and manage named maps. The same dialog is also used for the cluster density colour scale.



The colour map

The small squares below the colour map mark fixed colour points. Single click a small square to select it. Double-click a small square to change its colour. The colours between the fixed colour points are formed by linear interpolation. Click and drag the squares sideways to change their position. You cannot move the first and last fixed points. You can add a new fixed point by double-clicking on the colour scale away from any small square. The buttons above the colour scale are:

- Delete** This deletes the selected fixed point and is disabled if there is no selection.
- Insert** If there is a selected fixed point, this adds a new point next to the selected point on the side with the larger gap. If there is no selected point, this adds a new fixed point in the largest gap between existing fixed points. You can also add a fixed point by double-clicking the colour map clear of any fixed point. You can have a maximum of 10 fixed points.
- Even** This spaces the fixed points evenly across the colour bar.
- Reverse** This reverses the order of colours. Reversing twice restores the original.
- Complement** This replaces each colour by its complement, formed by subtracting each colour from white. Complementing twice restores the original.

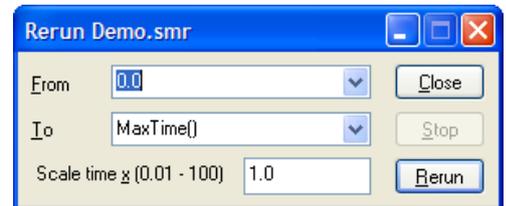
Name for this colour map You can save up to 10 user-defined maps in the registry. Set a name and click the **Add to Preset maps** button to add a new colour map to the list. If the list is full you must use the **Erase** button to remove a colour map before adding a new one. If you type a name that cannot be used, the **Add to Preset maps** button is hidden and a text message explaining the problem replaces it. Names of user-defined maps must be unique and may not include the vertical bar character.

The names are stored in the registry in the Spike2 area in the `Settings\ColourMap` folder with the names `ColourMap0` to `ColourMap19`. The map used for drawing sonograms is stored in `ColourMap0`, the map for clustering is `ColourMap1`. The user-defined maps are saved as `ColourMap10` to `ColourMap19`. The names `ColourMap2-9` are reserved.

Preset colour maps The drop-down list in this section holds several built-in maps with names that suggest how they are formed, for example **Rainbow**, **Thermal** and **Geography**. It also holds the names of maps that you have created and saved. Click **Select** to copy the map with the displayed name to the displayed colour map. The **Erase** button can be used to delete a user-defined map from the list and the registry. You cannot delete the built-in maps.

OK Click **OK** to accept the current map. It is saved in the system registry as `ColourMap0` for sonograms and as `ColourMap1` for clustering. The map is applied immediately.

ReRun This Time window option simulates data sampling for demonstration and training purposes. The **From** and **To** fields define the time range to use. The **Scale time** field sets the number of seconds of data to rerun every second. The **Rerun** button starts the simulation; the **Stop** button cancels it. Unlike real sampling, there is no sample control window and the output sequencer and on-line waveform replay are not available.



When the update point of the document reaches the right-hand edge of the window, the window scrolls to keep the update point at the right-hand edge of the screen. You can use the scroll bar to move through the document, but only up to the current end. If you move the scroll bar fully to the right, the window will scroll automatically, if the scroll bar is not at the right, the window will not scroll, but the scroll bar position will change as time passes to show its relative position in the document. The document will keep running until it reaches the end or you use the **Stop** button.

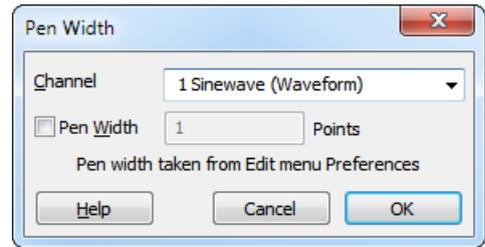
If there are open multimedia windows associated with the time window, these will also rerun and replay the audio signal to stay aligned with the current replay point.

In place of scrolled displays, you can have a paged update. Open the **View menu Display Trigger** dialog, select **None** as the trigger channel, check **Enable trigger** and click **OK**.

Link to Offline waveform output If you want to rerun a file to match a replaying waveform (either played through the 1401 DAC outputs or through a sound card in your computer), check the **Rerun the file to Match the waveform output** box in the **Offline waveform output** dialog.

Pen Width

The Pen Width dialog in a time or result view overrides the pen width set for data in all channels in the Edit menu Preferences dialog. The dialog is available from a context menu by right-clicking a channel, or from the View menu. Changes made with this dialog can be undone and recorded as a script. The script language equivalent is `ChanPenWidth()`. For XY views, use the XY Draw Mode dialog to set the pen width of a channel.

**Channel**

You can select a channel from the drop-down list or type in a channel or a list of channels. If more than one channel is set, the displayed information will be for the first channel in the list.

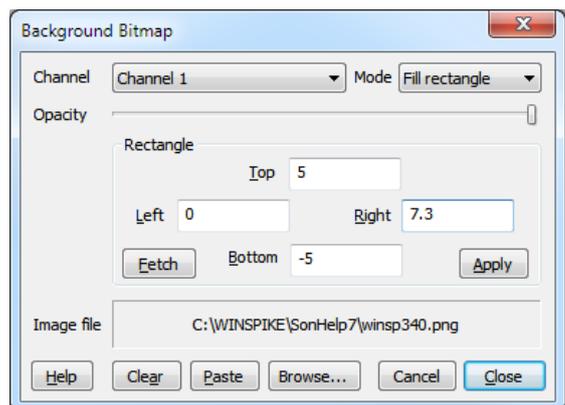
Pen Width

Check the box to set a pen width in points then click OK to apply it to all the channels in the list. If you clear the box, the standard pen width for data from the Edit menu Preferences tab will be applied on OK and you will not be able to edit the value. A point is 1/72 of an inch, which is about the same size as a pixel. If you set a size of 0, the pen size will be 1 pixel on all devices (which can be very thin on a printer). You can set a fractional point size, but the output pen will always be at least 1 pixel wide.

Channel Image

You can set a bitmap as the background to any channel in a Time, Result or XY view. This can be very useful in an XY view when the XY view is being used to map activity, for example, when the view is showing where spiking activity occurred in a maze, or when tracking a target. The menu command opens a dialog:

The dialog links a channel to a bitmap file on disk or to a clipboard image. Any channel in the view can be associated with an image. The name of the file appears in the Image file section of the dialog. The image is rendered on top of the channel background colour, but below all other items that are drawn. When channels are overdrawn, the image appears when the background colour of the channel would be drawn. The dialog fields are:

**Channel**

You can select any channel in the current view. However, if you are in an XY view, all channels share the background, so the choice of channel makes no difference. When you change channel, any Rectangle region changes that have not been applied will be lost and the current settings for the new channel will appear in the dialog.

Mode

There are three mode settings: No display, Fill Background and Fill Rectangle.

No display

No image is displayed. This allows you to have an image loaded, ready to display. This will save the time needed to load the image from a disk file.

Fill Background

The image is scaled so that it fills the available channel rectangle. You would use this when you want to display the entire image as a

background for the channel. Scrolling a time or result view in this mode look horrible; for this reason `Fill Background` is treated as `No Display` for Time views during sampling or rerun.

Fill Rectangle The image is scaled so that it fills a rectangle defined in x and y axis units. You would use this when the image must be aligned with the axes. For example, if the x and y axes represented co-ordinates in a maze, the image could be a picture of the maze. Scrolling works without any visual problem in this display mode.

Changes made to the mode field are applied immediately, so you can see the effect.

Opacity This field is a slider control that you can drag to the left to make the image transparent and to the right to make it opaque. Changes made by dragging are applied immediately.

Rectangle The fields within this area are disabled unless the mode field is set to `Fill Rectangle`. The `Left`, `Right`, `Top` and `Bottom` fields refer to the x axis and y axis co-ordinates of the image. In the current implementation, you cannot invert or reflect the image; if you do it will not appear.

Fetch This button sets the `Left`, `Right`, `Top` and `Bottom` fields to the values that would make the image fill the channel background.

Apply This button is enabled in `Fill Rectangle` mode if the current rectangle settings are legal numbers and do not match those of the channel. Click the button to apply the new values.

Browse... Click this button to open a file dialog in which you can browse for a suitable file. You can choose between windows bitmaps, JPEG images, PNG files and TIFF files.

Paste Click this to paste a clipboard image to the current channel. The Image file name will display as `<CB>`. Setting the clipboard as the source can have unexpected results as each time a saved window opens it will copy whatever image happens to be on the clipboard. Clipboard images are more useful to scripts, for example to copy a multimedia view video frame using `EditCopy()`, then setting the image with `ChanImage()`, avoiding the need to write the image to a disk file.

Clear Click this button to clear any image set for the channel. This releases any resources used to hold the bitmap.

Cancel Click `Cancel` to close the dialog and undo changes since the last channel change.

Close Close the dialog, leaving the associated channel in the current state.

This command is experimental (at version 7.06) and is implemented in a simplistic manner. If you reuse the same image for multiple channels we store the image multiple times, which could be a problem if you use huge images on many channels.

Show Gutter The gutter is an area to the left of the text in a text view, usually with a grey background, that can be used to select lines of text, except in a script view where clicking on it sets break points. The gutter is also used in script views to display the current line pointer. This menu command shows and hides the gutter, equivalent to the `Gutter()` script command. The gutter shares the same display background colour as line numbers. The gutter is normally visible.

Show Line Numbers You can choose to display line numbers in any text view. Line numbers usually have space for up to 5 digits, but if you need more than 99,999 lines of text you can use the `ViewLineNumbers()` script command to increase the displayed digits. This menu command makes the line number area visible if it is invisible and hides it if it is visible. You can change the appearance of the line number region with the View menu Font command. Select the Line number style, and a line number will appear in the example window so you can preview any changes.

Folding This menu item is available for Output sequencer and Script views to control the code folding options. The folding options are covered in detail in the Edit menu Preferences command in the General Tab for these view types. This View menu command gives you a quick way to change the folding style and to toggle all the folds. The folding styles are:

```
'Check view type
if ViewKind() then
    halt;
else
    Message("Time view");
endif;

'Type OK
const n% := 100, phi := 0.01;
var i%, x[n%], y[n%];
for i% := 0 to n%-1 do
    x[i%] := Sin(i%*phi);
    y[i%] := Cos(i%*phi);
next;
```

Arrows *Plus and minus* *Circles and Lines* *Squares and Lines*

The remaining folding commands are:

Toggle all folds This searches the text from the start to find the first fold, then sets all the folds in the file to the opposite state. The short-cut is `Ctrl+Shift+T`. To toggle an individual fold, click on the fold mark.

No folding This hides the fold margin. If this option is selected and text is folded, it is unfolded.

Spike Monitor The View menu Spike Monitor command opens a new window that displays all the WaveMark channels in the current time view in a grid (see the *Spike shapes* chapter for details). The command is disabled if there are no WaveMark channels in the time view.

Analysis menu

The **Analysis** menu is divided into several regions, all associated with the analysis of data or the processing of data into a different form. The first region holds commands that are associated with processing data from a time view into a result or XY view. The **New Result View** command opens a pop-up menu from which you can choose an analysis that generates a result window, for example a waveform average of one or more channels. The **Measurements** command is used to generate an XY graph from measurements based on cursor positions. The **Process** and **Process Settings** commands allow you to modify analyses created with the other two commands. The **Fit Data** command opens a dialog from which you can fit mathematical functions to sections of data in time, result or XY views.

The second region manages time view memory buffers, virtual channels and general channel operations. A memory buffer is like any other data channel in a time view except that it is held in computer memory and will disappear when the data file closes. It can hold any type of data that can be stored in a time view. Data can be added to memory buffers interactively, or you can import sections of other channels. You can save memory channels as part of the data document. Virtual channels are waveforms derived from combinations of other channels using the standard arithmetic operators add, subtract, multiply and divide. This section also holds commands to calibrate and process time view waveform and WaveMark channels, duplicate time and result view channels and save and delete time view data channels.

The third region holds commands to set the marker filter and process WaveMark data. The WaveMark commands are most useful when processing spike shapes, but they can be used for other tasks, such as detecting features (for example R waves in an ECG signal). The marker filter can also be used to isolate specific event markers, for example for conditional averaging.

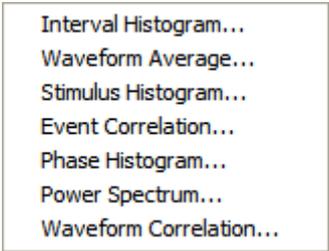
The last region holds the digital filter command. You can both generate and apply Finite Impulse Response digital filters to time view waveform channels to generate new memory channels or permanent channels in the data file.

New Result View



This command is available when the current window is a time view. It is a pop-up menu from which you can select an analysis type. This leads into a dialog where you define the parameters to construct a result window.

A result window holds arrays of data that can be drawn as histograms or as waveforms. There is one data array for each channel you analyse. Some types of analysis can store additional data, for example the peri-stimulus histogram can store event times for a raster display and the waveform average can store extra data so that you can display error bars in the result.



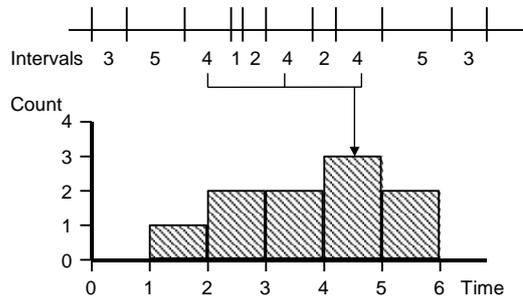
- Interval Histogram...
- Waveform Average...
- Stimulus Histogram...
- Event Correlation...
- Phase Histogram...
- Power Spectrum...
- Waveform Correlation...

Once you have set the required values in the dialog box, Spike2 creates a result window with all data values set to zero. A new dialog appears to prompt you to select a region of the original time window to analyse. The results of analysing different areas can be summed.

At the time of writing, the following analyses are implemented: interval histogram, waveform average, peri-stimulus time histogram, event time cross correlation, event phase histogram, power spectrum and waveform correlation.

Interval histogram

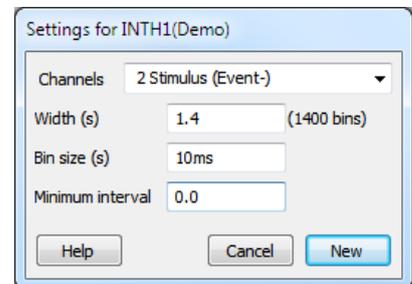
An interval histogram (INTH) displays the frequency histogram of the intervals between events on a selected channel. All times in a Spike2 data document are expressed as a multiple of the basic time unit set when the data document was created. Therefore you cannot get better time resolution than the basic time unit (usually in the range 2 to 50 μ s).



The picture illustrates how an interval histogram is formed. To make the diagram easy to understand, we have shown several events with all the intervals between them as a whole number of seconds. The interval histogram formed from these events has bins set with a width of one second each.

The interval between each pair of events is divided by the bin width to give the bin number to increment (fractional bin numbers are rounded down, the first bin is bin 0). In this case, the bin width is 1, so there is no rounding. In a more realistic case with an interval of 2.3 milliseconds and a bin width of 1 millisecond, bin number 2 (spanning the time period 2 to 3 milliseconds) would be incremented.

Event times are rounded down to the nearest base time unit, so an event time of t base units means that the actual time was greater than or equal to t and less than $t+1$ units. An interval of n units means that the real time interval between two events was greater than $n-1$ units and less than $n+1$ units. If you need to form interval histograms of very short periods, make sure that you have set the duration of the base time unit short enough to resolve the information you require.



The **Channel** field selects one or more event channels; type a list of channels or select from the drop down list. Each channel generates a histogram in the result view. If you type a channel list or use the selected channels, the order of the channels in the result view is the order of channels as you enter them.

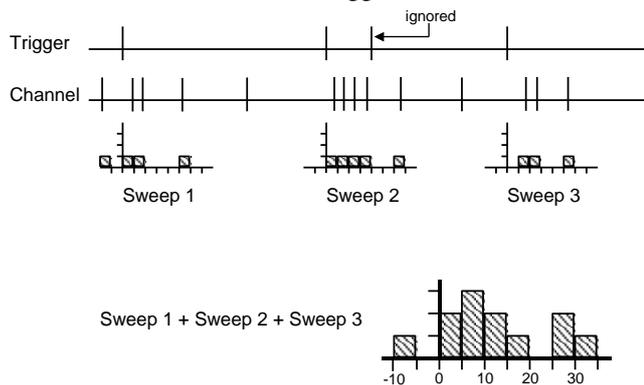
Width sets the time width of the histogram from left to right. The **Minimum interval** field is the smallest interval to put in the histogram and sets the left hand end of the x axis. It must be positive and is usually zero. **Bin size** is rounded to the nearest multiple of the underlying time units used in the time window. The number of bins ($\text{Width}/\text{Bin size}$) must be at least 1. The maximum size is limited only by available system memory.

The **New** button accepts the values in the dialog and creates a new result window. This dialog can also be activated by the **Process Settings** menu command, in which case the **New** button is replaced by a **Change** button. **Change** clears any previous results.

Click the **New** or **Change** button to open the **Process** dialog (see page 9-12) in which you set the time range of data to analyse. Only events that fall within the time range are considered, intervals that start or end outside the time range are ignored.

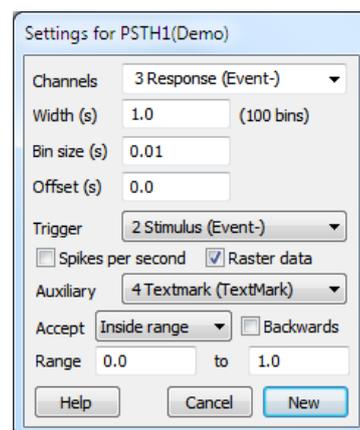
Peri-stimulus time histogram

A peri-stimulus time histogram (PSTH) forms a histogram of events on one channel around a stimulus event on a trigger channel.



Each trigger event that does not lie within the previous sweep generates a new sweep. Trigger events that fall in the previous sweep are ignored. For overlapped sweeps, use the event correlation analysis. Sweeps are accumulated to produce a histogram. You can scale the result into spikes per second rather than event counts.

In the Settings dialog, the Channel field selects the channels to analyse. The Width and the Bin size fields set the number of bins in the histogram. The Offset field sets the pre-trigger time to display. If the time offset is 0.5, the histogram time axis starts at -0.5. All times are rounded to the nearest multiple of the basic time unit.



The Trigger field sets a data channel to use as the trigger, or you can select no trigger. If you do not select a trigger channel, the process command uses the start time as the trigger time for a single sweep. If you select a trigger channel, the process command defines a time range to search for trigger events.

Spikes per second

If you check the Spikes per Second box, the histogram shows equivalent spike rate rather than the number of events in each bin. To change from event count to rate, Spike2 divides the number of spikes in each bin by the number of sweeps to give the spikes per sweep, and then divides the result by the bin width to give spikes per second.

You can change between count and spikes per second after you have created the result view and analysed data. Click on the result view, open the Process settings dialog and change the state of the Spikes per second check box, then click Change.

Raster data

If you check the Raster data box, the new result view stores the times of all the events that are added into histograms and you can to draw the result as a raster display. When you create the result view, Spike2 automatically duplicates each result channel and draws the duplicated channel in raster drawing mode above the histogram. As each event uses memory and takes time to draw, the maximum number of events is limited by the memory in your system and your patience. This should not be an issue unless you work with hundreds of thousands of spikes with the result view raster enabled.

Once you have created the result window and analysed data you can delete the raster data by opening the Process settings dialog and clearing the Raster data check box. You cannot add raster data to an existing result view without clearing all the histograms.

Auxiliary measurements

You can take one measurement per sweep from an auxiliary channel. If you choose a waveform or RealWave channel, the measurement is the channel value at the trigger point or 0 if there is no waveform data at the trigger. For all other channels types, the measurement is the time between the sweep trigger point and the first event on the

auxiliary channel after or before the trigger. Check the **Backwards** box to search before the trigger. Pre-trigger times are negative, times after it are positive. If no event is found, 0.0 seconds and the maximum time possible in the file are used as virtual event times.

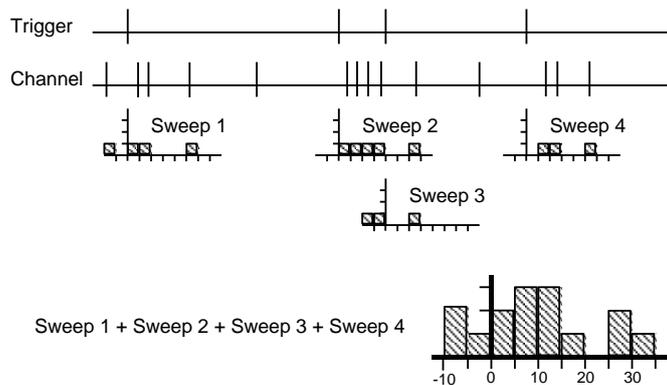
Accept sweeps You can accept all sweeps, or only those for which the measured value lies inside or outside the values in the **Range** fields. Sweeps with no data found are treated as outside the range. When the auxiliary channel holds waveform or RealWave data the units of the range fields are the channel units, for all other types the units are seconds.

Sorting rasters If you enable raster data, you can sort raster sweeps based on event latencies or waveform values. You set the sort mode from the Draw Mode dialog. Each raster sweep can store four sort values. The auxiliary measurements system uses the first sort value. Script users can access all the sort values with the `RasterSort()` script command.

Raster symbols Each raster sweep can store eight event times that you can choose to display as symbols with the Draw Mode dialog. If the auxiliary measurement is an event time, and an event was found, the time is saved as the first symbol time. Script users can access all eight symbol times with the `RasterSymbol()` command.

Event correlation

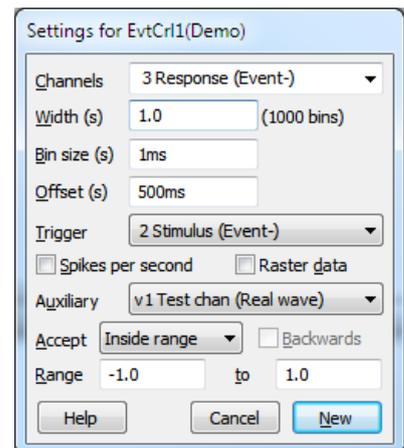
This command performs event cross-correlations and event auto-correlations between a trigger channel and one or more other event channels. A cross-correlation produces a measure of the likelihood of an event on one channel occurring at a time before or after an event on another channel. The result can be displayed as a spike rate or as a count.



The drawing illustrates the general principle of an event correlation. Every trigger generates one sweep of analysis (unlike a PSTH where a trigger event that falls within a sweep is ignored). The zero times of each sweep are aligned, and then the sweeps are accumulated to form a histogram.

The **Channel** field sets a channel or a channel list to analyse. The **Width** and the **Bin size** fields set the number of bins in the histogram. The **Offset** field sets the pre-trigger time to display, so if the time offset is **offset**, the histogram time axis starts at a time of **-offset**. All times are rounded to the nearest multiple of the basic time unit.

The **Trigger** field is a drop down list in which you can select a trigger channel or **Manual** mode. In **Manual** mode, the process command start time is the trigger time for one sweep. With a trigger channel, the process command sets a time range to search for trigger events.



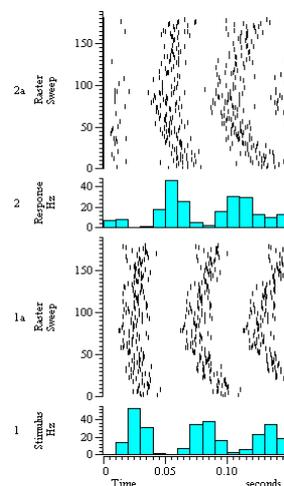
Spikes per second

If you check the **Spikes per Second** box, the histogram shows equivalent spike rate rather than the number of events in each bin. To change from event count to rate, Spike2

divides the number of spikes in each bin by the number of sweeps to give the spikes per sweep, and then divides the result by the bin width to give spikes per second.

Raster data If you check the **Raster data** box, the new result view stores the times of all the events that contributed to the histogram and you can use raster drawing mode for the result view channel. When you create the result view, Spike2 duplicates each channel and draws the duplicate in raster mode above the histogram.

As each event uses memory and takes time to draw, the maximum number of events is limited by the memory in your system and your patience. This should not be an issue unless you work with hundreds of thousands of spikes with raster enabled.



Auxiliary measurements You can take one measurement per sweep from an auxiliary channel. If you choose a waveform or RealWave channel, the measurement is the channel value at the trigger point or 0 if there is no waveform data at the trigger. For all other channels types, the measurement is the time between the sweep trigger point and the first event on the auxiliary channel after or before the trigger. Check the **Backwards** box to search before the trigger. Pre-trigger times are negative, times after it are positive. If no event is found, 0.0 seconds and the maximum time possible in the file are used as virtual event times.

Accept sweeps You can accept all sweeps, or only those for which the measured value lies inside or outside the values in the **Range** fields. Sweeps with no data found are treated as outside the range. When the auxiliary channel holds waveform or RealWave data the units of the range fields are the channel units, for all other types the units are seconds.

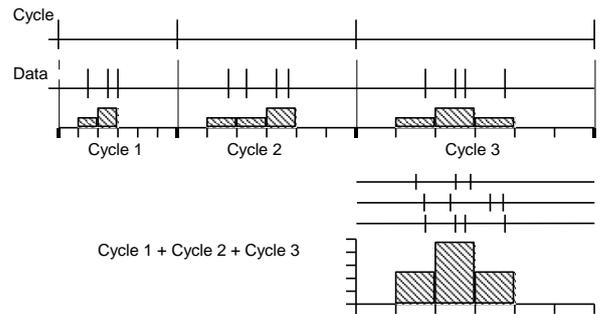
Sorting rasters If you enable raster data, you can sort raster sweeps based on event latencies or waveform values. You set the sort mode from the Draw Mode dialog. Each raster sweep can store four sort values. The auxiliary measurements system uses the first sort value. Script users can access all the sort values with the `RasterSort()` script command.

Raster symbols Each raster sweep can store eight event times that you can choose to display as symbols with the Draw Mode dialog. If the auxiliary measurement is an event time, and an event was found, the time is saved as the first symbol time. Script users can access all eight symbol times with the `RasterSymbol()` command.

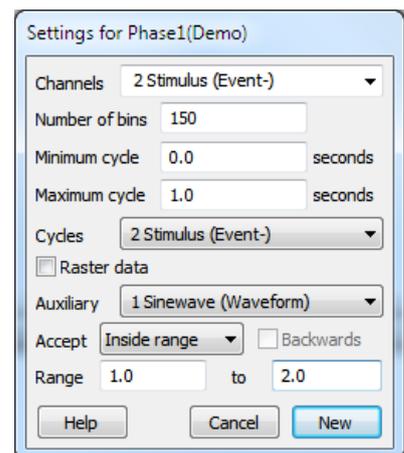
Auto-correlation For an auto-correlation, set the **Channel** and **Trigger** fields to the same channel. The analysis for an auto-correlation has one difference from that for a cross-correlation. In an auto-correlation, the correlation of each event with itself at time 0 in the histogram is ignored.

Phase histogram

A phase histogram is used to show how events are distributed with respect to a cyclical process that may vary in cycle time. One event channel marks the start and end of cycles. The end of one cycle is the start of the next. Events on another channel are placed in bins depending upon their position within each cycle.



The **Channels** field sets the channels to analyse. The **Cycle** channel sets the channel with cycle markers. The **Number of bins** field sets the bins per cycle. The width of each bin depends on the duration of each cycle. The **Minimum cycle time** and **Maximum cycle time** fields exclude cycles that are too long or too short to belong to the data and are set in seconds. Check the **Raster data** box to save the times of all events and duplicate each result channel in raster mode.



In the **Process** command, the start and end times determine the time range of the data in the **Cycle** channel to search for cycle markers. Cycles that are longer than the maximum time or shorter than the minimum time are ignored. If no **Cycle** channel is supplied, the start and end times are used as the start and end of a single cycle.

Auxiliary measurements

You can take one measurement per sweep from an auxiliary channel. If you choose a waveform or RealWave channel, the measurement is the channel value at the trigger point or 0 if there is no waveform data at the trigger. For all other channels types, the measurement is the cycle position of the first event on the auxiliary channel after or before the cycle start. Check the **Backwards** box to search before the trigger. Pre-trigger positions are negative. The cycle start position is 0 degrees, the cycle end position is 360 degrees. If an event is at time T_e and the cycle start and end times are C_s and C_e , the event position is $360 \cdot (T_e - C_s) / (C_e - C_s)$. If no event is found, 0.0 seconds and the maximum time possible in the file are used as virtual event times.

Accept sweeps

You can accept all sweeps, or only those for which the measured value lies inside or outside the values in the **Range** fields. Sweeps with no data found are treated as outside the range. When the auxiliary channel holds waveform or RealWave data the units of the range fields are the channel units, for all other types the units are degrees.

Sorting rasters

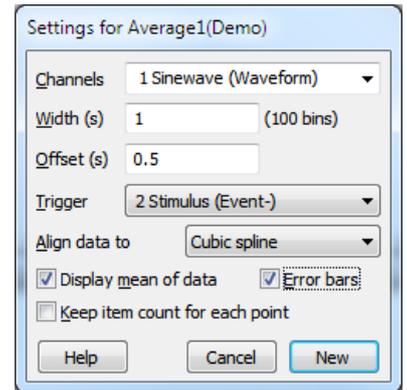
If you enable raster data, you can sort raster sweeps based on cycle positions or waveform values. You set the sort mode from the Draw Mode dialog. Each raster sweep can store four sort values. The auxiliary measurements system uses the first sort value. Script users can access all the sort values with the `RasterSort()` script command.

Raster symbols

Each raster sweep can store eight event times that you can choose to display as symbols with the Draw Mode dialog. If the auxiliary measurement is an event time, and an event was found, the time is saved as the first symbol time. Script users can access all eight symbol times with the `RasterSymbol()` command.

Waveform average

This command averages waveform channels with respect to a trigger signal, or if no trigger is set, averages waveform channels with respect to user-defined trigger times. The **Width** field sets the time range spanned by the average. The **Offset** field sets the start time of each sweep with respect to the trigger. The **Trigger** field sets the source channel for trigger times or **Manual** trigger mode.



Display mean of data

The **Display mean of data** chooses between a display of the mean of all the sweeps or a display of the sum of all the sweeps. If you return to this dialog after forming an average you can change this setting to redraw the result without needing to process the data again.

Error bars

If you check the **Error bars** box, extra information is saved with the result so that you can display the standard deviation and standard error of the mean of the resulting data. The **Channel Draw Mode** command controls the display of the error information.

Align data to trigger

In general, trigger times will fall between samples of the waveform channel. If you are looking at features that are only a few sample points in duration, for example transient responses, the data may change significantly between the sample points. You have a choice of four methods to set the alignment of the data to the trigger.

Next data point

The data is aligned at the next sample after the trigger time. This is the method used by all Spike2 versions before 6.01.

Nearest data point

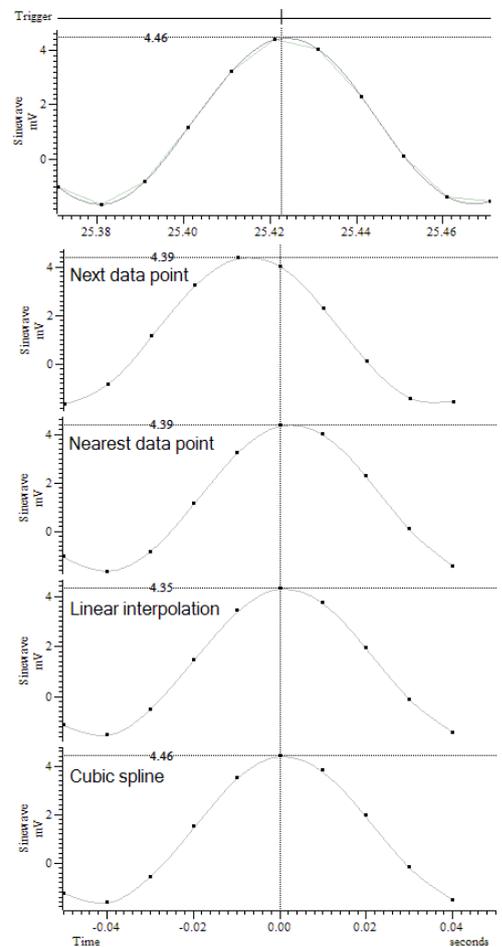
The data is aligned at the nearest sample to the trigger point. This is the best method to use if you don't want to interpolate.

Linear interpolation

The data values are aligned to the trigger by linearly interpolating between the data points.

Cubic spline

The data values are aligned to the trigger by cubic spline interpolation. This is the best method to use if your data contains no components at frequencies above half the sample rate. It is also the slowest method.



The picture shows the effect of each of the methods when accumulating a single sweep of data. The top trace shows the original data with the trigger point, overlaid with both linear and cubic-spline interpolations of the data. The four lower traces show what is added into the result for each of the methods.

The best method to use depends on your data. If your data is known to contain no frequencies above half the sample rate (which is what you would hope to be the case),

using a cubic spline will produce the best results. If this is too slow, you could try linear interpolation, which is faster to compute, but which will tend to reduce the amplitude of narrow features.

If your data contains transients with frequency content above half the sampling rate, cubic interpolation will become inaccurate and the best you can do is choose the *Nearest data point* method. The *Next data point* method is provided for backwards compatibility with older versions of Spike2.

Keep item count for each point

If you check this box, each point in the result view keeps track of the number of items that were averaged to create it. If you do not check the box, all points are presumed to have averaged the same number of items (the sweep count). The box makes a difference when sweeps of data are truncated or interrupted by gaps or the start or end of the channel data. If you do not check the box, missing data is presumed to hold zeros, which may cause discontinuities in the average. If you check the box, missing data makes no contribution to the average. If you use this option, saved result view files will not be compatible with versions of Spike2 before 5.16 and version 6.00.

New (Change) button

The New button (or Change when used from the Process Settings command) closes the dialog, creates a new result window and opens the Process dialog, described below. With a trigger channel, the events between the start and end times set in the process dialog are triggers. In Manual mode, the start time is the trigger for a single sweep.

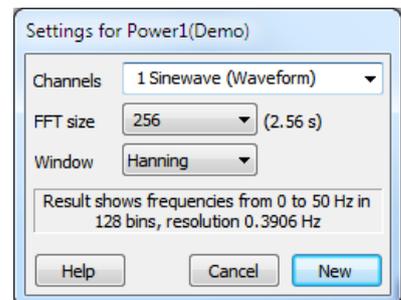
Sweep count

The result window for an average keeps track of the number of data sweeps that have been added into the average. A sweep is counted if a least one channel adds at least one data point into the average. It displays the mean waveform by accumulating the data and dividing by the number of sweeps or, if the *Keep item count...* box was checked, the item count for each point.

The start of the section of data to add to the average is found by subtracting the time in the *Offset* field from the trigger time. The data channel is then aligned to the trigger by the method set in the *Align data to trigger* field. If there is a gap in the waveform data, such that there is no data point that falls within time from the start of the section to the start plus the width of the average, no data is added.

Power spectrum

This command creates a result window that holds the power spectrum of a section, or sections of data. The result of the analysis is scaled to RMS power, so it can be converted to energy by multiplying by the time over which the transform was done. You can transform multiple channels, but they must have the same sample rate. Spike2 uses a Fast Fourier Transform (FFT) to convert the waveform data into a power spectrum.



The FFT is a mathematical device that transforms a block of data between a waveform and an equivalent representation as a set of cosine waves. The FFT that we use limits the size of the blocks to be transformed to a power of 2 points in the range 16 to 16384. You set the FFT block size from a drop down list. The result window ends up with half as many bins as the FFT block size. When you use the Process command, the selected area must hold at least as many points as the block size, otherwise no analysis is done.

The result window spans a frequency range from 0 to half the sampling rate of the waveform channel. The width of each bin is given by the waveform channel sampling rate divided by the FFT block size. Thus the resolution in frequency improves as you increase the block size. However, the resolution in time decreases as you increase the block size as the larger the block, the longer it lasts.

The mathematics behind the FFT assumes that the input waveform repeats cyclically. In most waveforms this is far from the case; if the block were spliced end to end there would be sharp discontinuities between the end of one block and the start of the next. Unless something is done to prevent it, these sharp discontinuities cause additional frequency components in the result.

Windowing of data

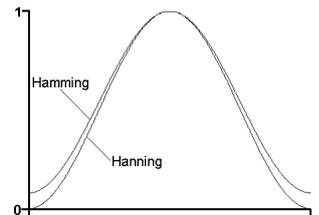
- No Window
- Hanning
- Hamming
- Kaiser 30 dB
- Kaiser 40 dB
- Kaiser 50 dB
- Kaiser 60 dB
- Kaiser 70 dB
- Kaiser 80 dB
- Kaiser 90 dB

The standard solution to this problem is to taper the start and end of each data block to zero so they join smoothly. This is known as *windowing* and the mathematical function used to taper the data is the *window function*. The use of a window function causes smearing of the data, and also loss of power in the result.

You can find all sorts of windows discussed in the literature, each with its own advantages and disadvantages; windows shaped to have the smallest side-lobes spread the peak out the most. By reducing the side-lobes you decrease the certainty of where any frequency peak actually is (or the ability to separate two peaks that are close together). Spike2 implements the following windows:

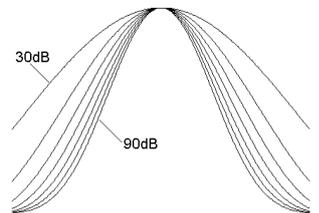
No Window Use this if there is one sine wave, or if more than one, they all have similar amplitude. This has the sharpest spectral peaks, but the worst side-lobes.

Hanning This is a good, general purpose, reasonable compromise window. However, it does throw away a lot of the signal. It is sometimes called a “raised cosine” and is zero at the ends. If you are unsure about which window would be best for your application, try this one first.



Hamming This preserves more of the original signal than a Hanning window, but at the price of unpleasant side-lobes.

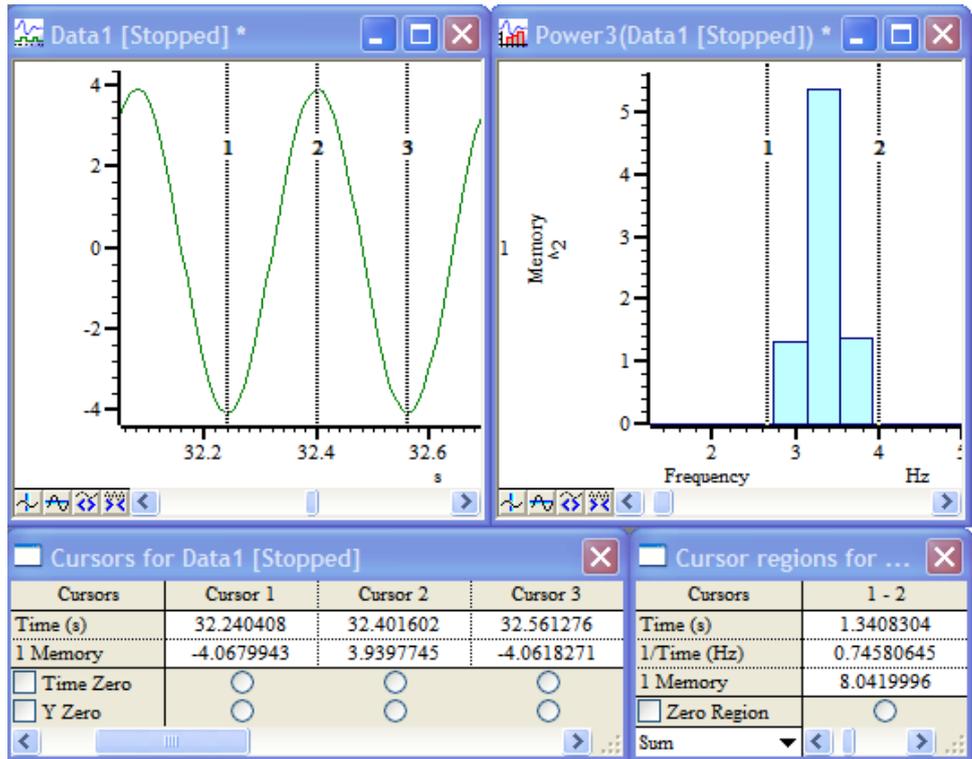
Kaiser These are a family of windows calculated to have known maximum side-lobe amplitude relative to the peak. Of course, the smaller the side-lobe, the more signal is lost and the wider the peak. We provide a range of windows with side-lobes that are from 30 to 90 dB less than the peak.



Power spectrum of a sine wave

If you sample a pure sine wave of amplitude 1 Volt and take the power spectrum, you will not get all the power in a single bin. You will find data spread over three (or more) bins, and the sum of the bins will be 0.5 Volt². The factor of 2 in the power is because we give the result as RMS (root mean square) power. This is illustrated by the example below where we have sampled a sine wave with amplitude 4.004Volts (peak to peak amplitude = 8.008, measured using cursors). We have formed the power spectrum of the signal using a 256 point transform and zoomed in around the bins where the result lies.

If the sampled data were a perfect sine wave we would predict a RMS power of 8.016 Volts² from this waveform (4.004² / 2). The cursor analysis of the power shows a total power of 8.042 Volts². This is about 0.45% above the predicted result.

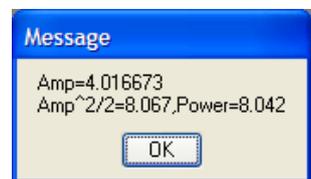


The predicted result is slightly low because the waveform samples used for the cursor measurements are unlikely to lie at the exact peak and trough of any particular cycle. Using the script language `Minmax()` function to find the maximum and minimum values over a wide time range gives a larger amplitude:

```

Evaluate a line of text:
var x,lo,hi;Minmax(1,2,40,lo,hi);x=(hi-lo)/2;Message("Amp=%f\nAmp^2/2=%3f,Power=%3f",x,x*w/2,8.042)
OK. 0.00 secs compile time. 514.31 secs execution time. 35 instructions, 498 globals, 0 locals, 1 strings.
    
```

You can find the Evaluate command in the Script menu if you want to try this. It gives a larger value for the amplitude (an over-estimate, as this is the difference of the highest and lowest extreme values) and now the power calculated from the amplitude exceeds the measured power by 0.3%. For an explanation of the text in the Evaluate window see *The Spike2 script language manual*.

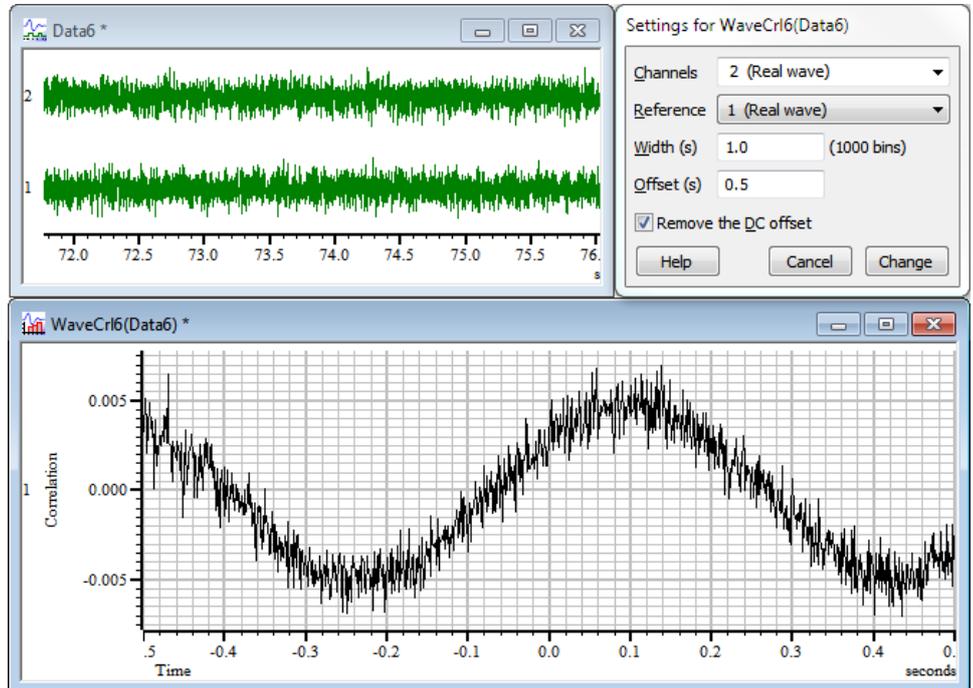


The duration of one cycle of the waveform (the time between cursor 1 and cursor 3) is approximately 0.32 seconds, a frequency of 3.125 Hz. Again, this is in agreement with the displayed power spectrum.

If you need access to the real and imaginary components of the FFT, you should consider using the `ArrFFT()` script language function.

Waveform correlation

This command creates a result view that holds the correlation between two waveform channels or the auto-correlation of a channel with itself. The correlation measures the similarity of two waveforms. The two waveforms must be sampled at the same rate. In the example below, the two waveforms hold random noise to which has been added a low amplitude 1.5 Hz sine wave. The sine wave in channel 2 is shifted forward by 0.1 seconds compared to channel 1.



The correlation is calculated by multiplying the two waveforms together, point by point, and summing the products. The sum is normalised to allow for waveform amplitudes and the number of points. This produces one result. The reference waveform moves one point to the right and the process is repeated to produce the next result. This is repeated for all the result bins. The results range between 1.0, meaning the waves are identical (except for amplitude) through 0 (un-correlated) to -1.0, meaning identical but inverted.

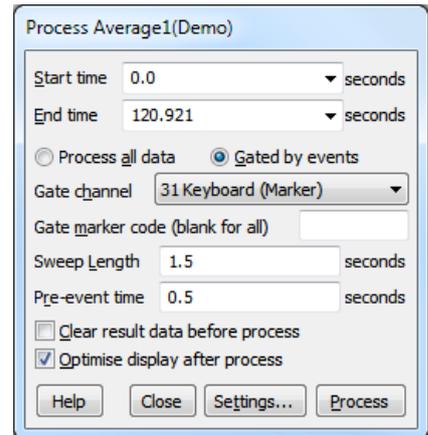
The bins in the output are the same width as the sampling interval of the two input waveforms. You can choose to remove the DC component from the signals before calculating the correlation. This can be important as DC offsets can dominate the result. The correlation is calculated in such a way that you can change the setting of the DC offset removal without recalculating the correlation. To do it, open the **Process Settings...** dialog again, click in the box and then click the **Change** button.

The **Process** dialog sets the time range of the reference waveform. The analysis is only done for regions of data in which both waveforms exist. If you calculate correlations over long sections of data, the calculations will take some considerable time! For this reason we do not recommend this as an on-line analysis (although there is nothing to stop you using it). The number of calculations (and hence the time taken) is proportional to the number of points in the result times the number of points in the reference waveform.

Process settings

This command opens the analysis setup dialog of the current result or XY window. The window must have been created by the **Analysis** menu **New Result** view or **Measurements** command. It is the same as the dialog that created the window except the **New** button is now a **Change** button. The **Change** button accepts changed settings and clears the result.

Process This command opens the Process dialog when the active window is a result or XY view attached to a time view. The dialog is also used to process measurements to a channel in a time view. It prompts you to select a time range of the data document to process. You can choose to add to the results of previous analyses, or you can clear the output. You can also choose to optimise the y axis of the result after the analysis is complete. The **Settings** button takes you back to the **Settings** dialog.



Start time and End time These fields set the time range to process. For triggered analyses, for example Waveform averages, peri-stimulus histograms and event correlations, the time range sets the trigger points to use. For non-triggered analysis, for example power spectra and measurements, the time range sets the data to analyse.

If a triggered analysis has **Manual** as the trigger channel, then the **Start time** is used as the trigger for a single sweep. In the special case of a Phase histogram with a **Manual** trigger, both start and end times are used to mark the beginning and end of a single cycle.

If there is more than one time window associated with the data document, then the drop down lists include the window number. For example, **View(-2).Cursor(1)**, meaning the position of cursor 1 in the second duplicate time window.

Gated analysis You will normally want to **Process all data** in the time range. However, you can also chose to process only those sections of the time window that lie within a user-defined time of particular events. For example, you might use events to mark a specific treatment type. You could then analyse data that fell within a time range of the treatment. To do this, select **Gated by events** and four extra fields are enabled.

Gate channel and Gate marker code You must choose an event, marker or marker-derived channel to define the gate times. If the channel holds markers, you can specify a marker code as a two digit hexadecimal code or as a single printing character. Only events in the time range that match the code will be used as gates. Leave the field empty to use all markers in the time range in the current channel marker filter mask.

Sweep length and Pre-event time These two fields define the length of time to process around each gate event and how far before each event to start processing. If gate events are closer than the sweep length, the sweeps are merged; the data is processed once only, not once per overlapped gate event.

Processing When you click **Process**, Spike2 evaluates the **start** and **end** fields and processes the data selected by the dialog. The dialog remains on screen until removed with **Close**. This allows you to accumulate the results of processing different areas. When processing to a time view channel that is not a memory channel, you can append data at the end of the channel only; you cannot add new data before the end of the channel.

Check the **Clear xxxx before process** box to clear all result data before the results of processing are added. The **xxxx** depends on the type of processing. Check the **Optimise Y axis after process** box to rescale the results to display all values.

Breaking out of Process Processing operations can take quite a time, especially in large data documents. If Spike2 detects a lengthy process operation, it displays a progress dialog in which you can cancel the operation. You can also stop processing early with the **ESC** key.

Process command with a new file

This menu command is available when the current window is a result or XY window attached to a sampling data file. It activates a modified version of the process dialog. This dialog also appears when you create a result or XY window while sampling. It is also used to control processing measurements to a channel during sampling.

This dialog gives you control over when and how the result is updated. You can select Automatic, Gated by events or Manual updates. The dialog contents depend on the update mode. The two check boxes operate in the same way as in the Process dialog, described above. The Settings button takes you to the Process Settings dialog, Close removes the dialog, Apply and OK both apply the settings, but OK also closes the dialog.

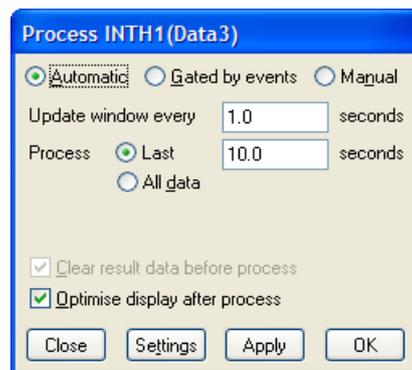
Automatic mode

In Automatic mode, the result updates as close to a user-defined interval as possible. Set the interval to 0 for the most frequent updates.

The Clear before process check box is grey as each of the two Process modes has only one useful setting. The modes are:

Last Re-calculates the result for the time period set using the most recent data. Use this mode to follow changes. Clear before process is always checked.

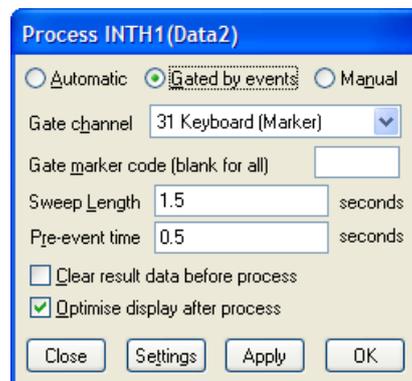
All The result of processing new data is added to the analysis. Clear before process is always unchecked.



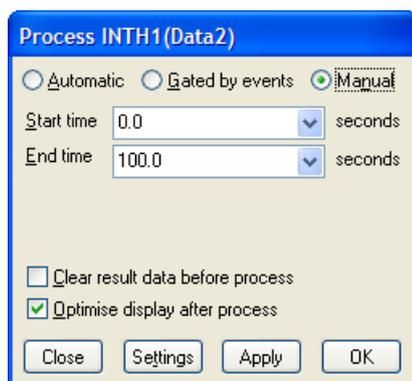
The Leeway for processing field is visible only when processing with active cursors to an XY view or time view channel. It sets the time to allow after cursor 0 for other cursors and measurements. If your analysis does not generate any data points, check that this value is large enough.

Gated by events mode

Gated by events mode analyses data around a specific event. You specify the sweep length, the start point of the analysis region relative to the trigger point, the channel to use as a trigger point, and if this channel is a marker, the marker code to cause processing (see the Sampling data chapter for details of marker codes). To accept all marker codes set a blank Gate marker code. If you do not want the results to accumulate for all analysis periods, check the Clear before process box in the Process dialog.



Manual mode

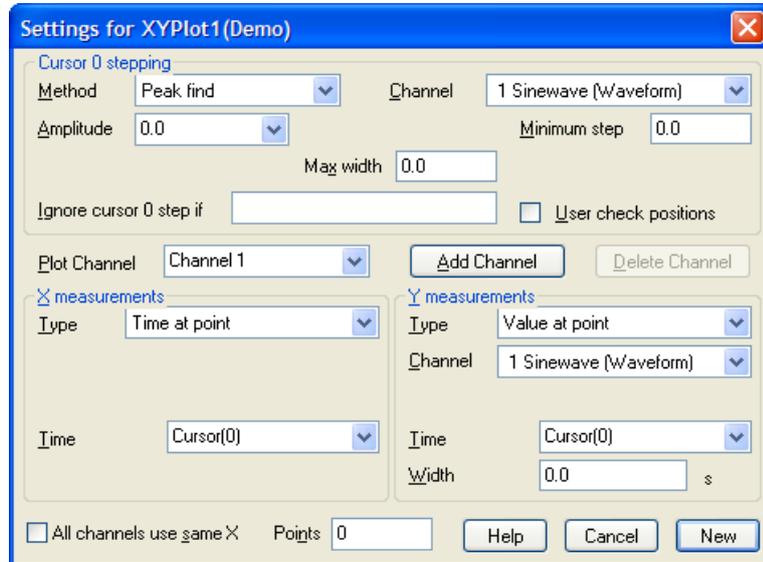


If you select Manual updates you must provide a start and end time for analysis. When you click Apply or OK, nothing will happen until the time range set for the start and end time is available in the data document.

At that point, the window contents will be calculated and the window will update. If you wish to process a different area, set a new time range and click the mouse on the Apply or OK button again.

Measurements to XY views

You can generate trend plots in an XY view both on-line in real time and off-line. Use the Measurements menu item and select XY view. This command generates a new XY window with one or more channels of data derived from the current time view.



The basic idea is that cursor 0 steps through the data following a user-defined rule. This can be as simple as *move to the last cursor 0 position plus 1 second* or it can be a complex data-searching algorithm. For each cursor 0 position, we take an x and a y measurement and pass the (x,y) point to a channel in the XY view. Of course, each time cursor 0 moves, all active cursors reposition themselves. Using this mechanism, a very wide variety of measurements can be taken and added to the result.

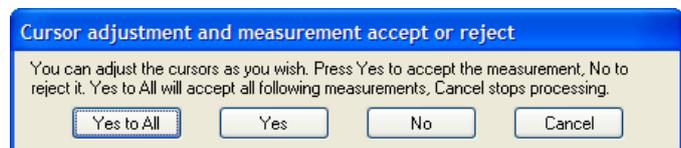
A PROCESS command sets the data region to search for cursor 0 step points, exactly as for the result view analysis commands. You can process data both on-line and off-line. Processing can take quite a long time if there are many cursor 0 steps to do. A progress dialog appears in a long process to give you the opportunity to cancel processing.

Cursor 0 stepping

This dialog area controls how cursor 0 moves through the data. If cursor 0 is already set to a suitable active cursor mode, this mode is copied to the dialog. You can set any active cursor mode that can iterate through data (see the *Cursor menu* chapter for details of all the active cursor modes). These are: Peak find, Trough find, Rising threshold, Falling threshold, Outside thresholds, Within thresholds, Slope peak, Slope trough, Turning point, Data points and Expression. The other fields in this section depend on the stepping mode.

You can use the Ignore cursor 0 step field to reject the result of a cursor step operation and step again. If this field is not blank and evaluates without error, and the result is true (not zero), the current cursor step is skipped. Expressions usually involve cursor values, for example $\text{Cursor}(0) > \text{Cursor}(1)$. You will normally leave this field empty.

If you check the User check positions box, you are given the opportunity to adjust the cursor positions after each step and before each measurement. This is ignored if the data file is being sampled.



Plot Channel

The trend plot can generate up to 32 channels of data in the XY view. The dialog starts with one channel. You can set the channel title in the Plot Channel box and create

additional channels with the Add Channel button. The Delete Channel button deletes the current channel; you cannot delete the last channel.

X and Y measurements These two areas have identical functionality. They generate the x and y values that are passed to the XY view for each point. The fields displayed in these areas depend on the measurement Type field. Most measurement types depend on times. For a useful result, you will set these to times relative to cursor 0, or to an active cursor that was positioned relative to cursor 0. The measurements are described in detail later in this chapter.

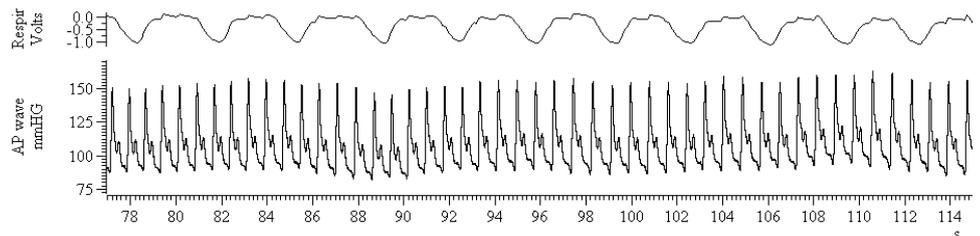
All channels use same X With multiple plot channels check this box to use the same x measurement for all channels. This is most commonly used when the x measurement is a time.

Points in plot If you set this field to a non-zero value, for example 20, the first 20 data points will appear, as you would expect. Subsequently, each added point causes the oldest point to be deleted. This is most useful on-line, for example to show pressure-volume curves or to display eye movements when you have horizontal and vertical input data.

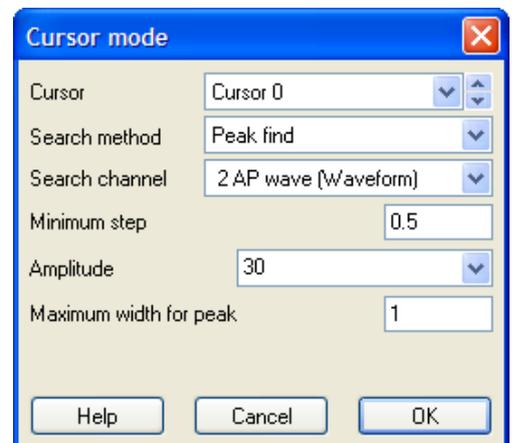
New and Cancel The New button creates the XY view and opens the Process dialog, ready to set a time range to step cursor 0 through to process the data. If you return to this dialog after creating the XY view, New becomes Change. Cancel closes the dialog.

Tabulated output You can tabulate the results in an XY view with the Edit menu Copy as Text command.

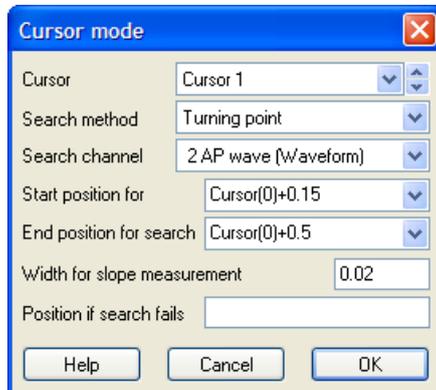
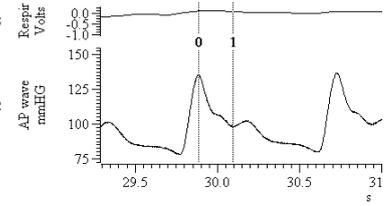
Example plot As an example of a trend plot, consider this data file containing an arterial blood pressure on channel 2 and a respiration signal on channel 3. Let us suppose that you are interested in the position of the dichotic notch (the small downward blip after the peak of the blood pressure) relative to the peak of the blood pressure.



The first step is to decide how to step through the data. The obvious method is to locate the blood pressure peaks. I used the Cursor menu Active modes... command to select cursor 0, set the search method to Peak find, the search channel to 2 and Amplitude to 30. Any Amplitude from 5 to 40 would work as each cycle is at least 40 mmHg from peak to trough and the biggest wobble in between is around 5 mmHg. I set the Minimum step to 0.5 and the Maximum width for peak to 1 to reject artefacts (although this data does not have any). I checked that this was working with Ctrl+Shift+right and cursor 0 stepped from peak to peak.



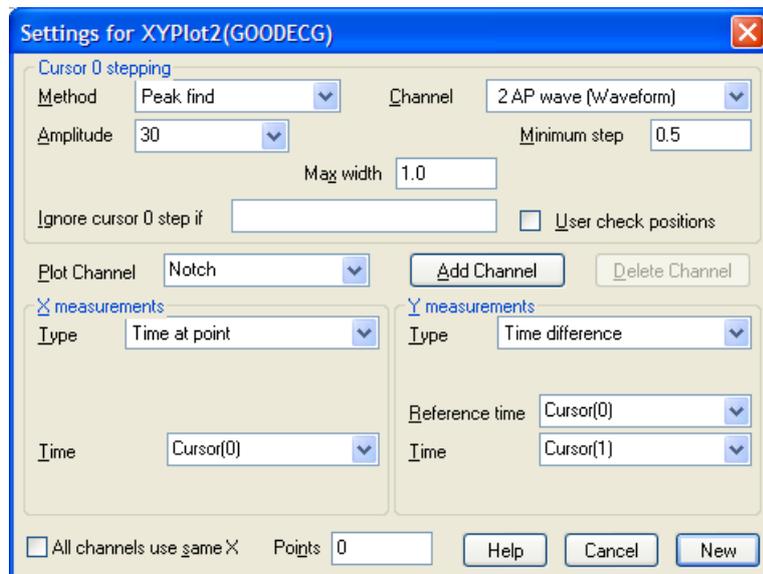
The next step is to locate the notch. I added cursor 1 to the window with `Ctrl+1`, and then opened the cursor Active modes dialog for cursor 1. This time I set the method to Turning point, set channel 2, set the search range to start at `Cursor(0)+0.15` and end at `Cursor(0)+0.5` and set the Width for slope measurement to 0.02 seconds.



I used Turning point mode and not Trough because the amplitude of the peak after the notch may be very small. I started the search just after the peak because the peak is also a turning point (where the slope changes sign) and I wanted to exclude it. From a visual inspection of the data I could see that the notch was always more than 150 ms past the peak, so by starting the search at `Cursor(0)+0.15` I avoided the possibility of detecting the change of slope between the peak and the notch.

The end of the search is set so that any reasonable notch will be found. The hardest item to set is the Width for slope measurement. This needs to be big enough so that noise in the signal doesn't cause false turning points, but small enough so we don't miss a small one. I checked it was working with `Ctrl+Shift+right` and now both cursors stepped along the data. Position if search fails is left blank so that if the turning point cannot be found, no measurement is taken.

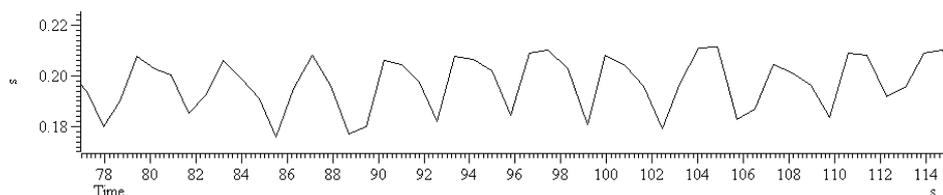
The final step is to generate the graph. I used the Analysis menu Measurements command and selected XY View. When the dialog opens it automatically picks up any active cursor setting from cursor 0.



I wanted to plot the notch position relative to the peak against the position of the peak. To do this I set the X measurement to be Time at point and selected `Cursor(0)` to be the point. I set the Y measurement to Time difference and set the Reference time to `Cursor(0)` and Time to `Cursor(1)` to form `Cursor(1)-Cursor(0)` as my measurement.

All that remained was to click New to create the XY window and open a Process dialog. I adjusted the time window so that it displayed the data I wanted to process, then

set the start and end times to `XLow()` and `XHigh()`, selected **Process all data** and clicked the **Process** button. The picture below shows the result. For each heart beat in the time range, we have a plot of the time delay from the peak to the dichotic notch.



Of course, we could have made a wide variety of measurements. The x axis need not be time, it could be a value derived from a different channel. For example, looking at the results, we can see that there seems to be a link between the notch position and the respiration signal, so we might have set the x axis to the value of the respiration channel at the cursor 0 time.

Measurement fields The X measurements and Y measurements areas determine the x and y values passed as a pair into each XY channel. The Type field sets the measurement method. The contents of the measurement area depend on the measurement type. The following fields are used:

Channel The data channel used for a measurement. Select the data channel from the list.

Time This is a time for use either as a value, or the time at which a measurement of the data channel is to be made. This field will usually be set to an expression that contains a cursor position, for example `Cursor(0)-1.3` or `(Cursor(0)+Cursor(1)/2)`. You can select expressions from the drop down list.

Expression This is used for Expression measurement mode only and is the expression to be used for a measurement. This will usually be a view expression that evaluates to a time.

Prompt This is used in User entered value measurement mode only, and is the prompt to use to request data values from the user.

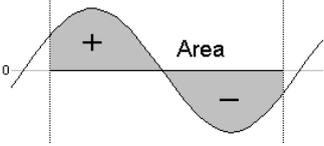
Reference time This field is used in the same way as the Time field. It specifies a time to subtract from the time in the Time field for the Time difference measurement, or the time at which to measure the data value to subtract for the Value difference measurement.

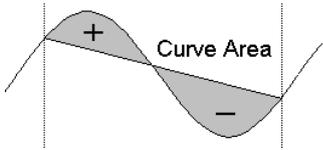
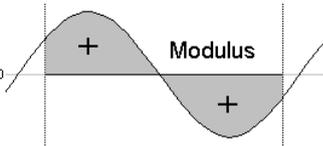
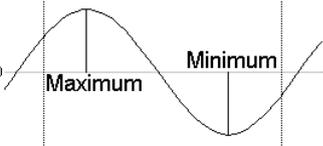
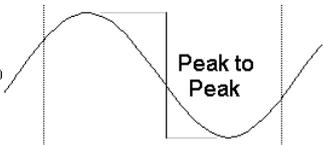
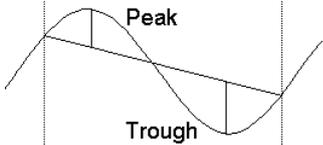
Width This field is used to set the width of a point measurement around a time. If you set this to 0, the nearest data point is used, otherwise a mean value over the time range from `Time-Width/2` to `Time+Width/2` is used.

Start time, End time These two fields always occur as a pair and identify a time range in which to take a measurement. Both fields will accept an expression for a time and you can select likely expressions from the drop down lists.

Minimum step Cursor 0 does not use the Start time and End time fields. Instead, you set the minimum step from the last cursor 0 position. The search range for cursor 0 ends at the end of the file (or the start if you are searching backwards) and starts at the minimum step away from the last cursor 0 position.

Measurement types In the description of each measurement type we refer to waveforms and events. By waveform, we mean either a waveform channel, or a WaveMark channel drawn as a waveform. The measurements are:

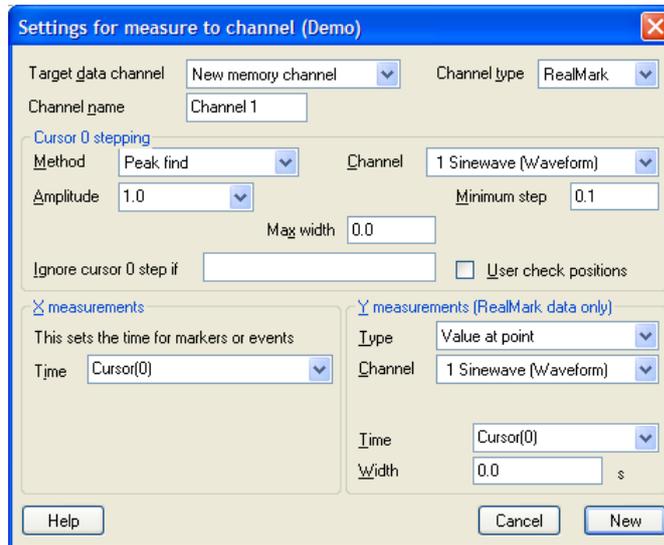
- Value at point** This is the value of the nominated Channel at the specified Time. If Width is non-zero, the measurement is the mean value over the time range from Time-Width/2 to Time+Width/2. For a waveform or a channel drawn as a rate, mean frequency or instantaneous frequency, the result is in y axis units. For all other event drawing modes, the result is the time of the first event at or after Time.
- Value difference** This is the value of the nominated Channel at Time minus the value at Reference time. If Width is non-zero, the measurement is the mean value from -Width/2 to Width/2 around each time point. For a waveform or a channel drawn as a rate, mean frequency or instantaneous frequency, the result is in y axis units. For all other event drawing modes, the result is the time difference between the first event at or after Time and the Reference time.
- Value above baseline** This is the same as Value difference, except that the Width measurement is applied only to the reference time value. The expectation is that the reference time is set in the middle of a baseline region, and value is a spot value at Time.
- Value ratio** This is identical to the Value difference measurement except that instead of subtracting the values, we divide the value of the nominated Channel at Time by the value at the Reference time. Attempted division by zero does not produce a measurement.
- Value product** This is identical to the Value difference measurement except that instead of subtracting the values, we multiply the value of the nominated Channel at Time by the value at the Reference time.
- Time at point** This is the value of the expression in the Time field. In most cases, this will be a cursor position, for example: `Cursor(0)`.
- Time difference** This is the value of the expression in the Time field minus the value of the expression in the Reference time field. In most cases, both these values will be cursor positions.
- Fit coefficients** A fit is requested on the channel set by the Channel field after each cursor iteration and the value of the coefficient selected by the Coefficient field is the measurement. For this to be useful, you must define a fit for the channel using iterating cursors for the start and end times. If a valid fit exists for the channel you can choose the coefficient by name, for example: a0: Amplitude, a1: Time constant, a2: Offset, otherwise you select the coefficient as a0, a1, a2 and so on.
- Expression** Type in an expression (usually one that involves cursor positions) and this is evaluated at each step to produce the value.
- Use entered value** You will be prompted to enter a value at each step of the cursor 0 position. If you enter a non-blank prompt, this is used when you are asked for a data value.
- Area** If the Channel is a waveform, the result of this measurement is the area between the waveform channel and the y axis zero over the time range set by Start time and End time. The area is positive for curve sections above zero and negative for sections below zero. For all other channel types this is the count of events in the time range.
- 
- Mean** For waveforms, the result is the mean value of the data points between Start time and End time. For other channel types this is the number of events in the time range divided by the time range. This could be thought of as the mean event rate in the time range.
- Slope** This has meaning only for a waveform. The result is the slope of the best-fit straight line by the least squares method to the data between Start time and End time.

<i>Sum</i>	For a waveform or a WaveMark channel drawn as a waveform this is the sum of the values of the data points. For all other channels it is the same as the Area measurement.	
<i>Curve area</i>	This measurement only has meaning for waveforms. It is the area between the straight line joining the ends of the waveform data in the time range and the data points. Curve areas above the line add to the result; areas below the line subtract from the result.	
<i>Modulus</i>	This has meaning only for waveforms. It is the same as the Area measurement except that all areas (above and below the zero line) count as positive. If you use this on an event channel no measurement is taken.	
<i>Maximum, Minimum, Extreme</i>	These have meaning only for a waveform. They measure the maximum and minimum values in the time range and Extreme measures the larger of the maximum and minimum value ignoring the sign. For example if the maximum value was 6 and the minimum value was -7, the extreme value would be 7.	
<i>Peak to Peak</i>	This measurement has a value for waveforms only. It is the difference between the maximum and minimum value of the channel in the time range Start time to End time .	
<i>RMS amp</i>	This measurement has a value for a waveform only. It is calculated by summing the square of each data point, dividing the sum by the number of data points and then taking the square root of the result.	
<i>SD</i>	The standard deviation has a value only for a waveform. It is calculated by finding the mean of the data, then summing the squares of the differences between each data point and the mean, dividing the sum by the number of data points minus 1, and finally taking the square root of the result.	
<i>Peak and Trough</i>	These measurements only have meaning for waveforms. These values are the maximum positive and negative distances between the waveform and a straight line joining the end points of the waveform in the time range Start time to End time . The Peak value is always greater than or equal to 0. The Trough value is always less than or equal to zero.	

Measurements to a data channel

This command generates a new event, marker or RealMark data channel in the current time view. The new channel holds the results of measurements made with the active cursors. You can use this command on-line as well as off-line. For example, if you were recording a blood pressure or ECG signal, you could use this to pick signal peaks to provide a real-time heart rate channel, directly from the blood pressure signal.

Use the **Measurements** menu item and select **Data channel** to open the dialog. You can also activate this dialog by right clicking on a channel that was created by a measurement and selecting **Process Settings** from the channel pop-up menu.



Target data channel When you are creating a new channel this field holds a drop-down list of possible channels. You can choose an unused disk channels or **New memory channel**. If you open the dialog after creating a channel, the field holds the channel number. You cannot replace an existing channel from this dialog; you must delete the existing channel first.

Channel type Choose from RealMark, Marker or Event data. If you change the channel type after processing this will delete all the channel data when you click **Change**. For RealMark data, the target channel has one real value attached to each marker. The value is set by the Y measurements field.

Channel name This field sets the name of the new channel. For RealMark data, the channel units are copied from the Channel used for the Y measurement.

Cursor 0 stepping The fields in this dialog area are identical to those in the **Measurements to XY** views.

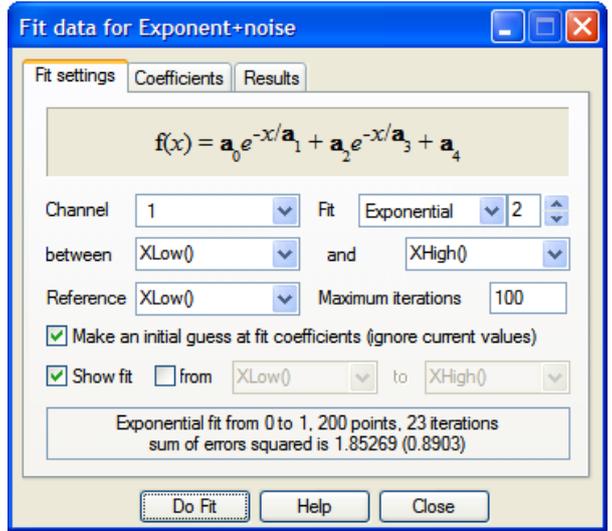
X measurements The **Time** field in this dialog section sets the time of each item that is added into the data channel. You will normally type in an expression that is related to cursor positions.

Y measurements The fields in this dialog area are identical to those in the **Measurements to XY** views. You can use this section of the dialog when the type of the target channel is RealMark otherwise this section is disabled. The value extracted here sets the RealMark data value.

New/Change This button is labelled **New** when you are creating a channel and **Change** if you return to this dialog after creating the channel. If you change the channel type and click **Change**, any previous data stored in the channel is deleted. You are asked to confirm that this is what you intended to do.

When you click the **New** or **Change** button, the **Process** dialog opens for you to choose the time range to process to generate data for the new channel. If you are sampling data into the time view, the on-line version of the **Process** dialog opens.

Fit Data This command opens a tabbed dialog from which you can fit mathematical functions to channels in a time, result or XY view. In a time view you can fit to an event-based channel as long as the data uses a display mode that has a y axis. If you fit data to a channel in a result view and error bars are displayed, the fit minimises the chi-squared value, otherwise the fit minimises the sum of squares of the errors between the data and the fitted curve. In addition to best-fit coefficients and an estimate of how much confidence to place in them,



you also get an estimate of how likely it is that the model you have fitted to your data can explain the size of the chi-squared value or sum of errors squared. If your data does not have error bars, these estimates are based on the assumption that all data points have the same, normally distributed error statistics. The dialog has three tabs:

- Fit settings** Set the fit type and range of data to fit and range to display
- Coefficients** Set the starting point for your fit and optionally fix coefficients
- Results** Display the fitting results and residual errors

The three buttons at the bottom of the dialog are common to all pages. The **Help** and **Close** buttons do what they say. **Do Fit** attempts to fit with the current fit settings.

Fit settings This page of the Fit Data dialog controls the type of fit, the data to fit and what to display. The area at the bottom of the window gives a synopsis of the current fit state. Fields are:

Channel You can select a single channel from the current view. If this is a time view, the channel must have a y axis. If you change the display mode of an event-based channel, any fit associated with the channel will most likely become invalid. In a result view, the fit is done to the basic channel data, even if the channel is displaying raster data.

Fit The fit to use is defined by its name and the order of the fit (a number). For example, an exponential fit allows single exponents or double exponents. The window at the top of the dialog displays the mathematical formula for the fitting function. The following fits are currently supported (N is the maximum order allowed):

Name	N	Comments
Exponential	2	This fit includes an offset. You can force a zero offset in the coefficients page. Set a local reference point for the fit, otherwise the even-numbered coefficients may become too large to be useful.
Polynomial	5	These fits do not require starting values for the coefficients.
Gaussian	2	If you attempt to fit two overlapping peaks you may need to manually adjust the guesses for the peak centres to get convergence.
Sine	1	You can fit a single sinusoid with an offset. If the frequency guess (in radians) is not reasonably close, the fit may not converge.
Sigmoid	1	This fits a single Boltzmann sigmoid by an iterative method.

Range You fit data over a defined x axis range, set by the **between** and **and** fields. You can choose values from the drop down list or type in simple expressions, for example `Cursor(1)+1`. There must be at least as many data points to fit as there are coefficients.

For example, to fit a double exponential, which has 5 coefficients, you need at least 5 points. Most fits will use many more points than coefficients.

Reference This is the x axis position to use as the zero value of x in the fitting function. The most common value for this would be the start point of the fit. However, in some cases you may want this to be elsewhere. For example, in exponential fitting, you may want to calculate the likely amplitude of a trace at some position. Making this position the reference point makes it easy to calculate the amplitude (it is the sum of the even-numbered coefficients).

Maximum iterations All fits except the polynomial are done by an iterative process. Each iteration attempts to improve the coefficient values. The iterating stops when improvements in the fit become insignificant, the iteration count is exceeded, the mathematics of the fitting process suggests that the fit is not going to improve or there is a mathematical problem. This field sets the maximum number of iterations to try before giving up.

Make an initial guess The iterative fits need a starting point. There are built-in guessing functions that usually generate a starting point near enough to the solution that the fitting process can converge. If you check this box, these guessing functions are used each time you click the Do Fit button. Otherwise, each fit starts with the current values.

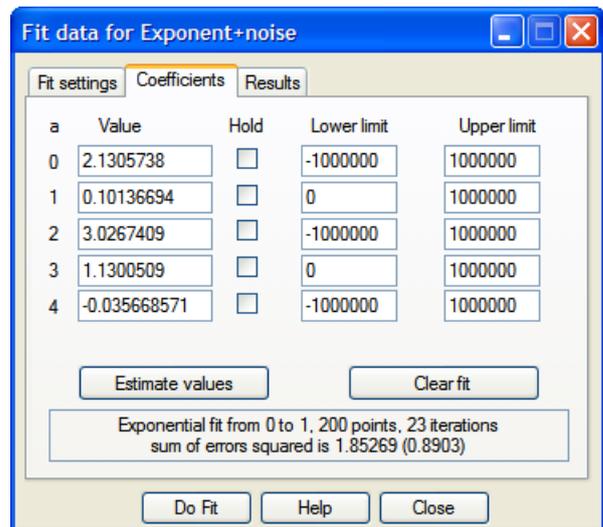
Show fit Check this box to display the current fit for the current channel. If the *From* box is checked, you can also choose the range over which to display the fitted data. If this box is not checked, the fit is displayed over the range that the data was fitted.

Coefficients This page of the Fit Data dialog lets you set the starting values for iterative fits. You can also use this page to hold some of the coefficients to fixed values and you can set the allowed range of values for fitting.

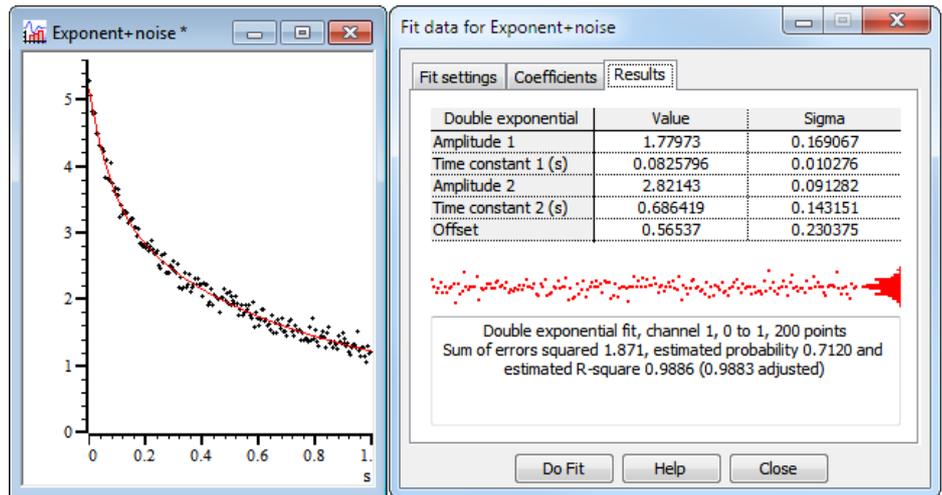
If you know the value of one or more of the coefficients, type the value in and check the Hold box next to it. For example, in an exponential fit you may know that the final coefficient (the offset) is zero.

The limit values are applied after each iteration. The fit may have to follow a convoluted path before it converges on a solution, so do not set the fit limits too close to an expected solution as this may prevent convergence.

The Estimate values button can be used to guess initial values for fitting based on the raw data. The Clear fit button removes the fit from the channel.



Results



Fitting to $y = 2.0 \cdot \exp(-x/0.1) + 3 \cdot \exp(-x/1.0) + \text{RandNorm}(.1, 0)$

The results page of the Fit Data dialog holds information about the last successful fit done with the dialog. The page has three regions: coefficient values at the top, a message area at the bottom, and a plot of the residuals (differences between the fit and the data) in the middle. The residuals are displayed immediately after a fit but will not be displayed if you close the dialog and reopen it.

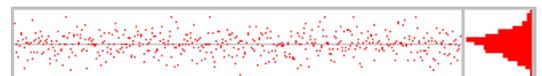
Coefficient values

The Value column holds the fitted value that minimised the chi-squared or sum of squares error for the fit. The Sigma column is an estimate of how the errors between the fitted curve and the original data translates into uncertainty in the fit coefficients given that the model fits the data and that the errors in the original data are normally distributed. If a coefficient is held, the Sigma value will be 0. The *Testing the fit* section gives more information on the derivation of these values and how to interpret them. You can select rows, columns or individual cells in this area, the use **Ctrl+C** to copy them to the clipboard. This also copies a bitmap image of the page to the clipboard.

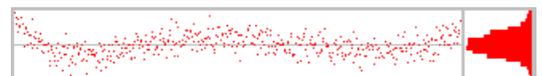
Residuals

This section of the page displays the differences between the data points and the fitted curve in the large rectangle and a histogram of the error distribution on the right. The error plot is self-scaling based on the distribution of errors; the plot extends from +3 at the top to -3 at the bottom times the RMS (root mean square) error. The grey line across the middle of the plot indicates an error of zero.

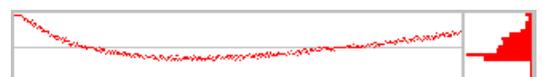
In the case that the data can be modelled by the fitting function plus normally-distributed noise, you would expect to see residuals distributed randomly around the 0 error line and the histogram on the right should resemble a normal curve.



If the data cannot be modelled in this way, you would expect to see evidence of this in the residuals. In this example (generated by fitting a cubic to data that was actually a double exponential), you can see that there are clear trends in the errors.



In extreme cases, the error due to the wrong model being used becomes much larger than the errors due to uncertainty in the data values, and you get a residual plot like this one.



Message area This area displays a summary of the fit information that you can select with the mouse and copy to the clipboard. The first line holds the type of the fit, the channel number, the ordinate range and the number of points in this range. For example: "Double exponential fit, channel 1, 0 to 50, 50 points".

The contents of the second line depend on the data source. If you are fitting a result view channel that has error information displayed, the second line displays the chi-squared error value for the fit and the probability that you would get a chi-squared value of at least that size if the function fits the data and the errors are normally distributed plus the R-squared value. For example: "Chi-square value 58.6, probability 0.5867, R-square 0.9886 (0.9883 adjusted)".

In all other cases, the second line displays the sum of the squares of the errors between the data and the fitted function and an estimate of the probability that you would get a sum of squares of errors of at least this size based on the assumption that the errors in the original data had a normal distribution that was the same for all points. For example: "Sum of errors squared 1.22, estimated probability 0.8553, estimated R-square 0.9886 (0.9883 adjusted)".

If the probability value is very low or very high, there are extra lines of information warning that the fitted function plus normally-distributed noise is unlikely to model the data, or that the errors in the original data have probably been over-estimated.

The R-square value is the standard "goodness of fit" value, which is the proportion of the variance of the data that can be explained by the fit, so a larger value is better. If you want to compare fitting different orders of polynomials or exponents you should use the adjusted values as these discount the improvement you would expect just because you increase the number of coefficient.

Context menu If you right click on this page you are offered a context menu that contains Copy, Log and Log Titles commands. The Copy command copies selected sections of the results, or all the results if there is no selection to the clipboard as text. It also copies the page as a bitmap. The Log command prints a one-line synopsis of the current fit to the log window. The Log Titles command copies a suitable set of titles for the logged data.

Testing the fit When you fit a model to measured data to obtain the best-fit coefficients, there are two questions you would like answered:

1. How well does this model fit the data? Put another way, how likely is it that this model plus some degree of random variation can explain my data set?
2. Given that the model does fit the data, how much confidence can I place in each of the fitted coefficient values?

When we talk about fitting curves to data, we are making the implicit assumption that you took measurements from some process that follows a model, and that this model can be expressed as a mathematical function with adjustable parameters, which are our fitting coefficients. Further, we assume that the measurements you make are not perfect; they have random variations with a known probability distribution about the correct value. To allow us to calculate likelihoods, we assume that this probability distribution is a normal (Gaussian) distribution. In the real world, or course, only some of this may apply. You may have no a priori knowledge of the distribution of errors in your original data, and this distribution may be anything but normal.

(Estimated) probability This value is an attempt to say how likely the model you have chosen plus normally distributed noise is to explain the measured data. If the residuals of your data show other than a random distribution around the zero line, the probability will be 0 (or close to it).

R-square (R^2 or $R2$ or r^2 value) This is a simple measure of *goodness of fit*. It describes what proportion of the variance of the data is explained by the model. Values range from 0 (the model accounts for none of the variation) to 1 (the model accounts for everything). A negative value is possible in non-linear fits if a horizontal line through the mean is a better fit than the model.

The adjusted R-square value is useful when you have a model that allows you to choose the order of the fit and you are not sure what order is appropriate. For example, when fitting polynomials you have a choice of order 1 through 5. Increasing the order improves the R-square value, even if all you are doing is fitting random variations in the data. The adjusted R-square value is less than the R-square value by an amount that compensates for the number of coefficients.

Chi-square fits In the ideal case, where you know the standard deviations of each data point, the fitting minimises the chi-squared value, which is the sum of the squares of the differences between the model and the data points divided by the standard deviation of data point values. Given a chi-squared value and the number of points it was measured from, we can calculate that probability of getting a chi-squared value at least this large, due to random variations in the data. This is the value given in the **Results** tab. Ideally, you would like to see a value around 0.5, meaning that you were equally likely to get a larger value as a smaller one. Values very close to 1 mean that, given the errors in each point, the data is too close to the model. Either the error estimates are too large, or the data has been "improved". Although you can hope for probabilities in the range 0.1 to 0.9, values down to 0.01 may occur for acceptable fits, and even smaller values can occur if your error distribution is not as normal as you thought.

Very low fit probabilities will occur if your data contains variations that are significant compared to the errors in



the input values and that are not included in the model. For example, if you are fitting exponents to a sampled waveform that includes perceptible mains interference, you can get a good fit (by eye) to the exponential data, but with a probability of 0.0000 as far as the mathematics is concerned because the model does not include the mains hum and cannot explain why the chi-squared value is so high.

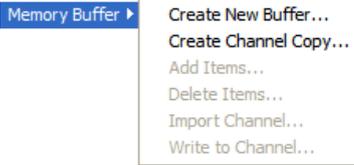
If we assume that the model fits the data, we can make an estimate of the standard deviation of the fitted coefficients. This means, that if we re-ran the experiment many times and fitted the data to each set of results, what would be the likely variation in the fitted coefficients. This is presented as the **Sigma** value in the results tab.

Least-square fits If there is no error information for each point, we assume that all the points have the same, normal error distribution and the fit minimises the sum of squares of errors between the model and the data. Because there is no independent estimate of the likely spread of the errors in the original data, strictly speaking, there is no way to give a probability of getting an error of at least this size.

However, we can say (though statisticians may shudder), "*Given that the model does fit the data, and that the errors all have the same, normal distribution, then the differences between consecutive errors should also be normally distributed with twice the variance of the errors*". We use this to estimate the standard deviation of the data and then we apply the probability test. We label this as *estimated probability*. The same comments about likely values apply as for the Chi-square fits, except that very small values may just mean that our estimation process fails for your data.

The coefficient **Sigma** values are calculated on the assumption that the model fits the data, that all the original points have the same standard deviation, and that the standard deviation of the original data can be deduced from the residual sum of squares errors.

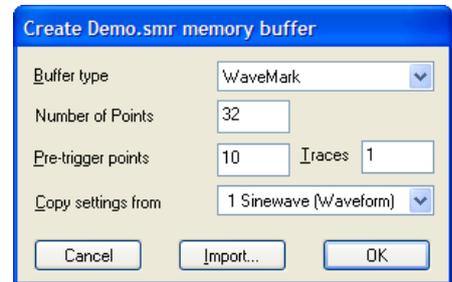
Memory buffer



Each data file has up to 300 memory channels. They hold data copied from existing channels, derived from waveform channels or entered by hand. You can display the memory buffers and use them like any other channels. The memory buffers can also be written to the data document. If you do not write them, the memory channels are lost when you close the file. Each memory buffer expands as events are added. The size is limited by available memory.

Create New Buffer

This command creates a new memory buffer channel of any type. The channel numbers are in the range 401 to 700. Previous versions of Spike2 used channels 101 to 200 and the numbers may change again in future versions. The memory channel numbers are displayed as m1 to m300. A new channel gets the lowest available number. You can use the Analysis menu Delete Channel command to remove memory buffers. OK creates the channel, Import creates it and opens the Import channel dialog.



Channels created with this command are not permanent. The data is kept in memory and is lost when the file is closed. If you need to make the data permanent, you must write it to a permanent channel with the Write to Channel option.

The fields in the dialog change depending on the type of buffer you choose to create. The example shown is for a WaveMark channel. The fields are:

Buffer type You can create a channel of any type: Event-, Event+, Level, Waveform, RealWave, Marker, TextMark, RealMark or WaveMark.

Number of ... This field is present for TextMark data where it sets the maximum number of characters to store with each mark, RealMark data where its sets the number of real values to store with each mark and for WaveMark data where it sets the number of waveform points to store with each mark.

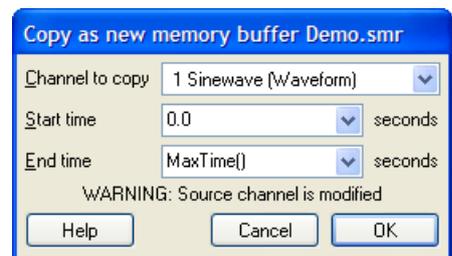
Initial level This field is present for Level event channels and sets the initial level as Low or High.

Pre-trigger points and traces These fields are present for WaveMark data only and set the number of waveform points before the trigger and the number of traces in each WaveMark item in the range 1 to 4.

Copy settings from When you create a waveform, RealWave or WaveMark buffer, Spike2 needs to know the spacing of the waveform points and the calibration settings. You use this field to indicate a channel that holds waveform, RealWave or WaveMark data and the buffer is given the same sampling rate and calibration. If you create a buffer from the script language you can choose a sample rate and calibration without reference to another channel.

Create Channel Copy

This command creates a new memory channel that is a copy of an existing channel. The source channel can be of any type. If the source channel has a channel process attached, or if it is a marker channel with an active marker filter, the new channel will contain modified data. You can also activate this command by right-clicking on a channel.

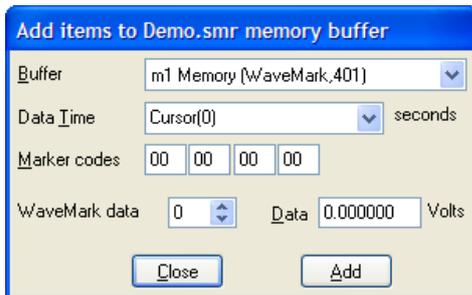


Channel to copy The source channel to copy. Copied items include the channel title, units and comment.

Start time, End time These fields let you choose a time range of the original data to copy to the new channel.

Add items to memory buffer

This command adds a data item to a memory buffer of any type. Click Add to place or replace data in the buffer at the specified time. Waveform and RealWave data is aligned in time with existing data



The fields depend on the channel type:

Buffer The memory buffer channel to use.

Event time The time at which to add data. If you have cursors enabled in the time view you can add data at the time of the cursor (as in the example). This field is present for all channels.

Marker These four fields are present for all Marker derived data types. You can set the marker codes appropriate for your data here.

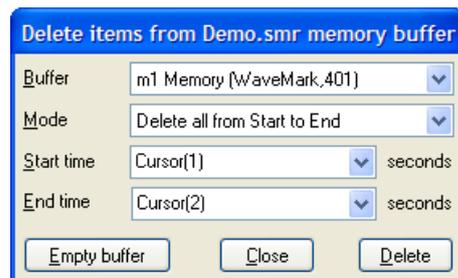
WaveMark/RealMark data This field is present for WaveMark and RealMark channels. Enter the index of the data point that you wish to set in the Data field. There is also a Trace field for WaveMark buffers with multiple traces.

Wave value This field is present for waveform and RealWave channels. You can type in a number or an expression that can include horizontal cursor values.

TextMark string This field is present for TextMark data only. Enter the text to appear in the new data item.

Delete items from memory buffer

This command opens a dialog in which you can delete one or more data items or a time range from a memory buffer. The Delete button removes one or more items as set in the dialog. The Empty buffer button deletes all data for this channel (it does not delete the channel). Close removes the dialog. The MemDeleteTime() script command has the same functions.



Buffer The memory buffer channel to use.

Mode There are four modes: Delete nearest to Start within Range, Delete all round Start within range, Delete first in range Start to End, Delete all from Start to End. The first two modes delete one or more data items around a time, the other two modes delete the first or all data items in a time range.

Start time The time range is Start time – Time range to Start time + Time range in the first two modes, and Start time to End time for the last two modes.

End time The end of the time range. This field appears in the second two modes.

Time range The time range around the Start time to search for data. This field appears in the first two modes. This field is usually set to a small value so that you can delete events close to the position of a cursor.

Import channel You can import data into a memory buffer selected by the **Buffer** field from a channel set by the **Channel** field. The **Start time** and **End time** fields set the time region to import data from. The **Mode** field is present if the source is a waveform channel and the destination is not, and sets the method to extract event times from the waveform.

To import a waveform channel to a waveform or WaveMark channel the ADC sampling rates must match. Apart from this restriction, you can move data from any channel to any buffer. During data import, values that cannot be extracted from the source are set to 0. For example, when importing an event channel into a WaveMark channel, there are no marker codes or waveforms in the source, so these are set to 0.

Waveform channels There are four modes to extract events from waveforms: **Peaks**, **Troughs**, **Data rising through level** and **Data falling through level**. Detected events are added to the memory buffer as markers with codes 2 (Peaks), 3 (Troughs), 4 (rising data) or 5 (falling data).

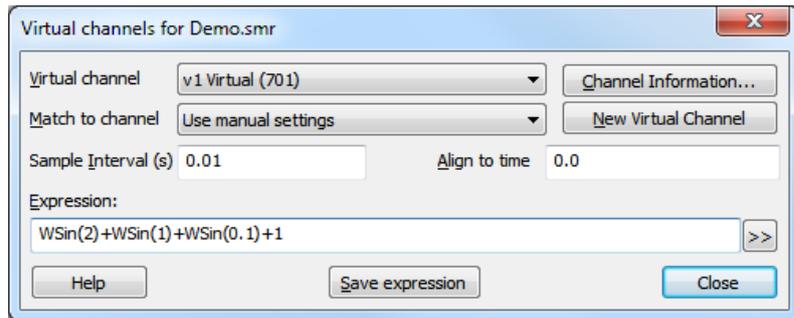
These modes use the **Minimum interval** field as the minimum separation of detected events, to filter out events caused by noise in the input waveform. Set this to 0.0 if you do not want to set a minimum period between detected events. The peak search mode looks for a peak followed by a fall of at least the **Size** field. The trough search mode looks for a minimum, followed by a rise of at least the **Size** field. The rising and falling level modes detect events when the signal crosses the **Level** in the selected direction.

Write to channel You can write a memory buffer to the data document. The **Type** field sets the format for the saved data. You can also **Append** the data to an existing channel. In **Append** mode, the data in the memory buffer must all occur later than data in the target channel. For modes other than **Append**, if you select a channel number that is already in use, you will be warned.

This command saves the entire contents of the memory buffer, regardless of any marker filter or channel process that may be attached to the channel. This is different from the **Save channel** command, which writes the data as displayed, including the effects of marker filters and channel processes.

Virtual Channels Virtual channels hold RealWave data derived by a user-supplied expression from waveform, event and RealWave channels and built-in function generators. No data is stored; the channels are calculated each time you use them. You can match the sample interval and data alignment to an existing channel, or type in your own settings. Channel sample intervals and alignments are matched by cubic splining the source waveforms, linear or cubic interpolation of RealMark data and by smoothing event rates. The script language equivalent of this command is `VirtualChan()`. You can save a virtual channel to disk (to remove the calculation time penalty) with the **Analysis** menu **Save Channel** command.

There are menu commands to create a new virtual channel or edit the settings of existing virtual channels. Both commands open the **Virtual channel** dialog:



- Virtual channel** Use this field to select a channel when you have more than one virtual channel.
- Match to channel** You can select an existing waveform-based channel (but not a virtual channel) from which to copy the sample interval and data alignment. Alternatively, you can select **Use manual settings** and type in the interval and alignment yourself.
- Sample Interval** This field holds the sample interval between data points in the virtual channel in seconds. You can edit the sample interval if you select **Manual Settings** in the **Match to channel** field. This field accepts expressions; for example, to set 27 Hz you can type $1/27$.
- Align to time** This field sets the time of a data point in the virtual channel. The time of any point and the sample interval completely defines all the sample times for the channel. You can edit the alignment if you select **Manual Settings** in the **Match to channel** field.
- Expression** This field holds an expression that defines the virtual channel. Expressions are composed of scalars, vectors and operators. A scalar is a number, such as 4.6 or $\text{Sqrt}(2)$. A vector is a list of data values that are derived from a channel or generated by a waveform function.
- Arithmetic operators* You can use the four standard arithmetic operators plus (+), minus (-), multiply (*) and divide (/) and comparison operators together with numbers, round brackets and some mathematic and channel functions. The result of combining a vector and scalar with an operator is a vector, for example, the expression $\text{Ch}(1)+1$ is a vector, being the data points of channel 1 with 1.0 added to each of them.
- Dividing by a scalar value of zero is an error. Dividing by a vector holding zeros is not an error and generates special floating-point numbers for positive and negative infinities. These can cause problems in subsequent calculations (and they are difficult to display).
- Comparison operators* You can also use comparison operators $<$, $<=$, $=$, $<>$, $>=$ and $>$ and there result of these is 1 if the comparison is true and 0 if it is false. In the expression $a \text{ op } b$, where a and b are vectors and/or scalars and op is a comparison operator, if either a or b is a vector, the result is a vector with value 1 at points where the comparison is true and 0 where it is false. For example $\text{Ch}(1) > 1$ has the value 1.0 where channel 1 is greater than 1 and 0 elsewhere. $\text{Ch}(1) <> \text{Ch}(2)$ is 1 wherever channels 1 and 2 are not the same. Be cautious using $<>$ (not equals) and $=$ (equals) as we are dealing with floating point numbers and exact equality may be compromised by arithmetic rounding.
- Operator precedence* Multiply and divide have a higher precedence than all the other operators which all have the same precedence. You can use round brackets to force other evaluation orders. Apart from that, evaluation is from left to right, so $a > b * c + d$ is interpreted as $(a > (b * c)) + d$.
- Channel functions* These functions take a channel of data (or a lower-numbered virtual channel) and create a vector holding data at the sample interval and alignment set for the virtual channel.

$Ch(n)$	n is a waveform, RealWave, WaveMark (first trace) or Level event channel to copy. Level event channels copy as 1 for high and 0 for low. Duplicate channels are allowed, but not duplicates of virtual channels.
$If(n, g)$	n is an event channel to convert to a waveform by linear interpolation of the instantaneous frequency. g is the maximum gap to interpolate across, in seconds. Omit g or set it to 0 for no limit to the gap.
$Ifc(n, g)$	The same as $If()$ except that cubic spline interpolation is used.
$Rm(n, g, i)$	n is a RealMark channel to convert to a waveform by linear interpolation of the data values, g is the maximum gap to interpolate over in seconds, and i is the data item to interpolate (the first item is 0). Set g to 0 for no maximum gap. You can omit i or i and g ; omitted values are set to 0.
$Rmc(n, g, i)$	The same as $Rm()$ except that cubic spline interpolation is used.

Spectral functions The following functions calculate the power spectrum of an input channel and generate a waveform illustrating how it changes with time. These functions are particularly useful in EEG analysis. The power spectrum is calculated using the Fast Fourier Transform (FFT), applied many times to span the data. Arguments common to all commands are:

- n The input waveform or RealWave channel. The maximum frequency for use in bands available to any of the commands is half the sample rate of this channel. Missing data values (gaps) in the input channel are treated as zero values.
- f The desired frequency resolution of the output. For a resolution of f Hz, the FFT must span at least $1/f$ seconds of input data using a power of 2 points, which limits the available resolutions; f is rounded down to the nearest available value. If the virtual channel sample rate is more than 8 times the frequency resolution we calculate the output for 8 times the resolution and interpolate the result. If you set f to 0, the FFT size is set to 256 points.
- l The low edge of a frequency range, in Hz. If optional and you omit it, 0 is used.
- h The high edge of a frequency range, in Hz. If this is an optional argument and you omit it, half the sampling rate of channel n is used.

These functions can take some time to calculate. If you find that these operations are slow, consider saving the virtual channel as a real channel so that you need only suffer the slow operation once. The available functions are:

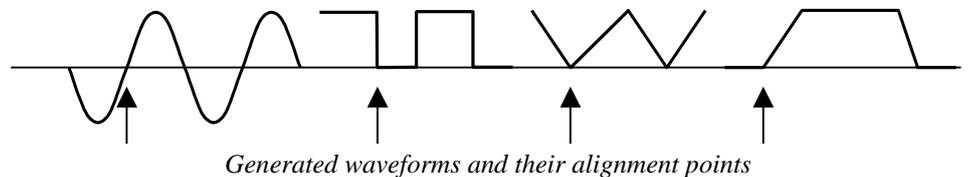
$Pw(n, f, l, h)$	Calculate the power in a frequency band from l to h Hz.
$Pw(n, f, l, h, r\{, s\})$	Calculate the power in the frequency band from l to h Hz divided by the power in the band from r to s Hz. If s is omitted, the band extends from r to half the sample rate of the input channel.
$SpE(n, f, p\{, l\{, h\}\})$	Calculate the frequency below which there is p percent of the power in the band from l to h Hz (Spectral Edge).
$MF(n, f\{, l\{, h\}\})$	Calculate the mean frequency in the band from l to h Hz.
$DF(n, f\{, s\{, l\{, h\}\}\})$	Calculate the Dominant Frequency. We search the band from l to h Hz for the frequency region with the maximum power. s sets the distance in Hz to smooth the data either size of each

Event kernel functions The following functions convert event channel n to a waveform by replacing each event by a kernel of unit area with a width set by w to the left and r to the right. You may omit r , in which case the kernels are symmetrical (r is set to w). You can set one (but not both) of r or w to 0 for a single sided kernel.

- Ec (n, w{, r}) This vector expression converts the event channel into a waveform by counting the events from -w to +r seconds around each bin.
- Et (n, w{, r}) This converts an event channel to a waveform by weighting the events from -w to +r seconds around each bin with a triangle function.
- Es (n, w{, r}) This converts an event channel to a waveform by weighting the events from -w to +r seconds around each bin with a sinusoid.
- Eg (n, w{, r}) This converts an event channel to a waveform by weighting the events from -w to +r seconds around each bin with a Gaussian function with a sigma (standard deviation) of w/4.
- Ee (n, w{, r}) This sets an exponential kernel. The time constant of the exponential is w to the left of each event and r to the right. The kernel extends to 8 times the time constant in each direction.

Waveform generation

These functions generate waveforms without the need for an input channel and produce output from time 0 to the end of the file. All the arguments are in seconds, except the sine wave frequency, which is in Hz. All these waveforms have unit amplitude. Use the standard maths operators */+ and - to scale and shift the output to different values.



If you attempt to create a cyclic waveform with a frequency above half the channel sample rate, no output will be generated. The a argument sets the time alignment; if omitted, the value 0 is used. You can generate the following outputs:

- WSin (f{, a}) Sine wave of frequency f Hz aligned so that phase 0 (the point where the rising sinusoid crosses 0) is at time a seconds. The amplitude of the output runs from -1.0 to +1.0. The sinusoid is most accurate at the start of a file; the accuracy falls off as the number of cycles increases (not usually noticeable).
- WSqu (l, h{, a}) Square wave with low period l seconds and high period h seconds aligned so that a low period starts at time a seconds. Both l and h must be greater than 0 seconds. The output level of the low section is 0, the output level of the high section is 1.
- WTri (r, f{, a}) Triangle wave with a rise time of r seconds and a fall time of f seconds aligned so that a rise starts at time a seconds. Either of r or f may be zero, but not both. The triangle output level is from 0 to 1.
- WEnv (r, h, f, a) Envelope with a rise time of r seconds, a hold time of h seconds, a fall time of f seconds with the rise starting at time a seconds. The output waveform is 0 before time a and after time a+r+h+f and is 1 during the hold time. At least one of r, h or f must be non-zero.
- WPoly (f, t, r, L) Polynomial in time from f to t seconds and 0 before f and from t. The value at time t is a function of (t-r) where r is a reference time. L is a list of 1 to 6 coefficients. WPoly (f, t, r, a, b, c, d, e, f) is:

$$y(t) = a + b*(t-r) + c*(t-r)^2 + d*(t-r)^3 + e*(t-r)^4 + f*(t-r)^5$$
- WT (s, e) Ramp from time s up to time e of value at time t of (t-s) with value 0 before s and from e. Omit e to run to the end of the file, omit both e and s to start from the beginning.

Mathematical functions	The mathematical functions mostly have the form <code>Func(x)</code> , where <code>x</code> can be either a scalar or a vector and <code>Func()</code> is the operation. If <code>x</code> is a vector, the result is a vector with the function applied to each value otherwise the result is a scalar.
<code>Abs(x)</code>	The absolute value of <code>x</code> (negative values are replaced by positive values of the same size). Use this to rectify a vector.
<code>Hwr(x)</code>	Half wave rectify <code>x</code> . Negative values are replaced by zeros.
<code>Sqr(x)</code>	This calculates the square of <code>x</code> . <code>Sqr(x)</code> is the same as <code>x*x</code> , but is faster, particularly when <code>x</code> is a vector.
<code>Cub(x)</code>	This calculates the cube of <code>x</code> . <code>Cub(x)</code> is the same as <code>x*x*x</code> but is much faster, particularly when <code>x</code> is a vector.
<code>Sqrt(x)</code>	This calculates the square root of <code>x</code> . When <code>x</code> is a vector, negative values are set to 0. If <code>x</code> is a scalar, negative values cause an error.
<code>Ln(x)</code>	Natural logarithm of <code>x</code> . Negative or zero <code>x</code> values generate -100 when <code>x</code> is a vector and an error when <code>x</code> is a number.
<code>Max(x, y)</code>	The maximum of <code>x</code> and <code>y</code> . This can be used to set a lower limit, for example <code>Max(Ch(1), 1)</code> is a vector with minimum value 1.
<code>Min(x, y)</code>	The minimum of <code>x</code> and <code>y</code> . This can be used to set an upper limit, for example <code>Min(Ch(1), Ch(2))</code> can be thought of as the value of channel 1 with the maximum value limited by the value of channel 2 (or vice versa).
<code>Exp(x)</code>	Exponential of <code>x</code> . If the result overflows, this is an error when <code>x</code> is a number and sets the largest allowed result when <code>x</code> is an array.
<code>Sin(x)</code>	Calculates the sine of <code>x</code> (<code>x</code> is in radians).
<code>Cos(x)</code>	Calculates the cosine of <code>x</code> (<code>x</code> is in radians).
<code>Tan(x)</code>	Calculates the tangent of <code>x</code> (<code>x</code> is in radians). The result can be infinite.
<code>ATan(x)</code>	The arc tangent of <code>x</code> . The result is in radians in the range $-\pi/2$ to $\pi/2$.
<code>ATan(s, c)</code>	The arc tangent of <code>s/c</code> in the range $-\pi$ to π using the signs of <code>s</code> and <code>c</code> .
<code>Poly(x, L)</code>	Replace <code>x</code> with a polynomial in <code>x</code> of order 1 to 5; <code>L</code> is a list of 1 to 6 coefficients. Use this to apply a non-linear calibration to a signal. For example: <code>Poly(x, a, b, c, d) = a + b*x + c*x*x + d*x*x*x</code>

You can insert spaces between operators and numbers and between round brackets and the items within them. You may not insert spaces between a function name and the opening bracket that follows it.

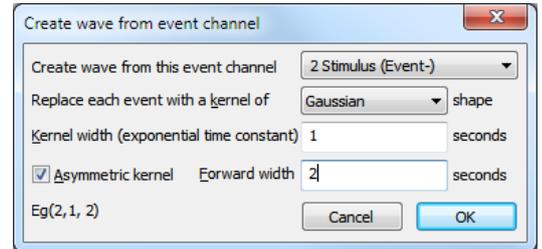
Example expressions If channels 1 to 3 hold waveform or RealWave data, then `Ch(2) - 2*Ch(1)` displays the difference between channel 2 and twice channel 1.

`Sqrt(Sqr(Ch(1)) + Sqr(Ch(2)) + Sqr(Ch(3)))` displays the square root of the sum of squares of three channels. You could use this to display the magnitude of the resultant of three perpendicular forces or movements. `Sqr(Ch(1))` is the same as `Ch(1)*Ch(1)`, but `Sqr()` is faster. To generate a polynomial function of the input it is much quicker to use `Poly()` than to use `Ch()`, `Sqr()`, `Cub()` and so on to generate a power series.

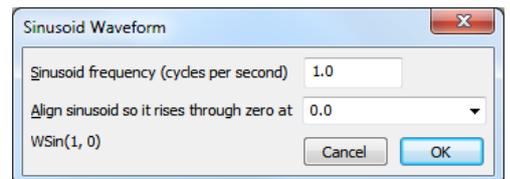
`((Ch(1) > (Ch(2) - .1)) * (Ch(1) < (Ch(2) + .1)))` displays the value 1 where channel 1 lies within 1 unit of the value of channel 2, and zero elsewhere. You might use an expression like this where channel 1 is the attempt of a subject to track motion described by channel 2.

Build expression There is no need to remember all the expression functions; just click the >> button and choose from a list of possible items to add. You can choose from:

Waveform from channel Select this to generate the commands that create a waveform from an existing channel. You build the commands using a dialog. All dialogs display the equivalent command in their lower left corner. The result replaces the selection in the Expression field. The example dialog illustrates creating a waveform from an event channel using a Gaussian kernel.

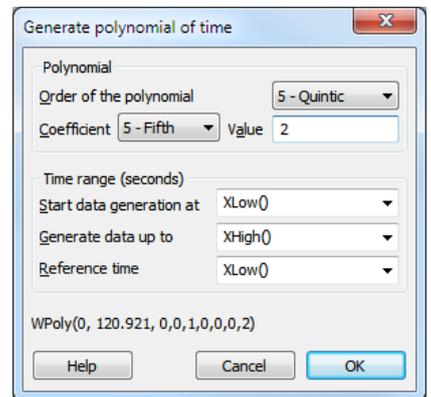


Generate waveform Select this to generate the commands that create a channel based on a sine, square or triangle wave or on a waveform envelope or based on linear time or on a polynomial of time. You build the commands using a dialog. All dialogs display the equivalent command in their lower left corner. The result replaces the selection in the Expression field.



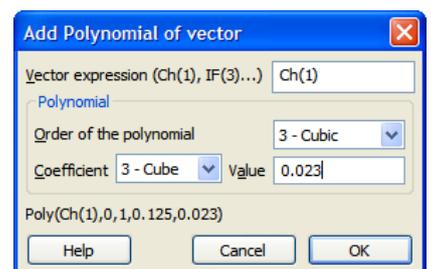
The `WPoly()` command is more complex, and can be used to generate polynomials in time relative to a reference time. Such curves can be used to create complex envelopes, or to subtract out a curve generated by the interactive curve fitting routines.

The `WT()` command generates a ramp from a start time up to, but not including an end time. The data is zero outside this time range. Within the time range, the ramp value is the current time minus the start time.



Mathematical functions The commands in this section apply a nominated mathematical function to the selection in the Expression field, which will usually be vector expressions. For example, if the Expression field holds `Ch(1) + Ch(2)` and you want to rectify the channel 1 data before adding it, select `Ch(1)`, click the >> button, select `Rectify` and `Absolute value`, then select `Rectify` the selection.

Mathematical functions: Poly() Select this to generate a polynomial. The *Vector expression* field is set to whatever was selected when this dialog was opened, or is set to `Ch(1)` if nothing is selected. This field is not tested for validity. Each element x of the vector is replaced by a polynomial in x . You can set the order of the polynomial (order means the highest power of x used) in the range 1 to 5. The command is:



$$\text{Poly}(x, a_0, a_1, a_2, a_3, a_4, a_5)$$

where x is the vector expression (you can also use a scalar, but this is not very useful) and the a_0 to a_5 are the coefficients of the polynomial. This expression generates the quintic:

$$y = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + a_4 x^4 + a_5 x^5$$

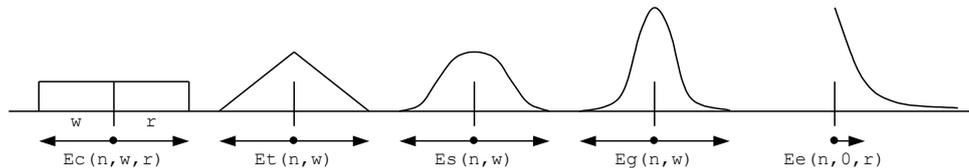
For lower order polynomials, omit the coefficients from the right. For example, for a quartic (fourth order polynomial), omit a_5 , for a cubic omit a_5 and a_4 . To set coefficient values in the dialog, use the *Coefficient* field to select the required coefficient and then set its value.

Mathematical operators These commands replace the selection with $+$, $-$, $*$, $/$ and with the comparison operators to remind you that you can use these operators to add, subtract, multiply, divide and compare vectors and scalars.

Previous virtual channel expressions This option lists expressions that you have used previously. Valid expressions are added to the list when you click the **Save expression** button, when you change to a different virtual channel and when you close the dialog with the **Close** button. The expressions are stored in the system registry. The most recently used expression is at the top of the list.

The Event to Waveform functions

The $E_C(n, w, r)$, $E_T(n, w, r)$, $E_S(n, w, r)$, $E_G(n, w, r)$ and $E_E(n, w, r)$ functions convert event channel n into a virtual wave and can be used anywhere in an expression that you can use the $Ch()$ function. They replace each event by a kernel (shape), centred on the event time. For an event at time t , the kernel extends from $t-w$ to $t+r$ seconds except for $E_E()$, which extends from $t-8w$ to $t+8r$ seconds. Normally you will omit r , in which case the shapes are symmetrical with r set equal to w . The resulting waveform is the sum of the kernels for all the events. The area of each kernel is unity, so the area under the waveform between any two times is the number of events in that time interval.



The $E_C()$ function simply counts the events in the time range. This is the fastest analysis method, but produces the most jagged output. The $E_T()$ function weights the events with a triangle function. This is slower than $E_C()$, but much faster than $E_S()$ and $E_G()$. The $E_S()$ function weights the events with a raised cosine. The $E_G()$ function weights the events with a Gaussian curve extending to 4 sigmas.

The $E_E()$ function is rather different from the others as it is not suitable for use as a smoothing function and is more likely to be used in a single-sided form. It weights the events with an exponent $\exp(-t/r)$ to the right and $\exp(-t/w)$ to the left (t stands for the time difference between the point and the time). The exponent extends to $8w$ to the left and to $8r$ to the right.

If none of these conversions are suitable, you can define your own weighting functions and create a new (not virtual) channel with the `EventToWaveform()` script command.

Duplicate Channels

This duplicates selected channels in the current time or result view. Duplicate channels share data, channel scales and comment with the parent and inherit the channel settings. Once you have duplicated a channel you can change title, display mode and y axis range independently of the original. With time view marker data, you can change the marker filter. The channel number of a duplicated channel is displayed as the original channel number plus a letter. The first duplicate of a channel gets the letter **a**, the second **b** and so on up to **z**, then **A** to **Z** are used. Duplicate channels are deleted with the **Analysis** menu **Delete Channel** command. New duplicates get the lowest available free letter.

Save channel

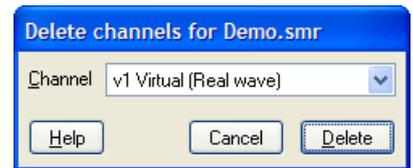
This copies the channel selected with the **Source channel** field to an unused channel on disk set by the **Destination channel** field. The channel is saved as displayed; any changes due to channel processes or marker filters are preserved in the output. The dialog displays a warning if this is the case. This command is not the same as the memory buffer **Write to channel** command, which saves the unmodified data in the memory buffers and allows you to set the destination channel type.



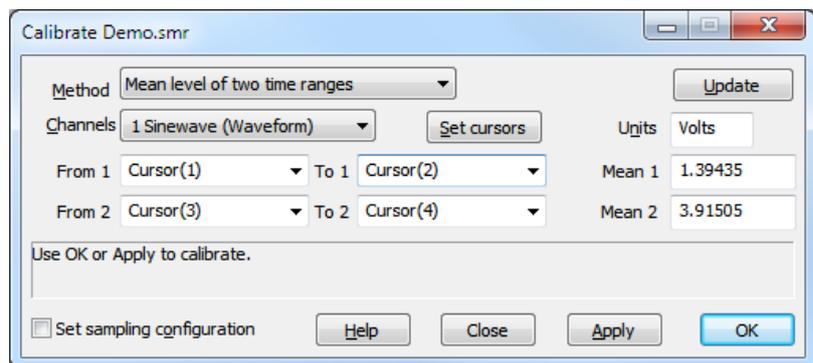
The `ChanSave()` script command also allows you to save channels, but gives you more control, including the ability to time shift data and to write data to a different file.

Delete channel

This removes a channel from a Spike2 data file permanently or deletes a channel created with **New Buffer** from the **Memory Buffer** command. If you delete a channel with duplicates, the duplicates are also deleted. Deleting a duplicated channel removes the duplicate only. You can also delete channels (except the last one) from XY views and duplicated result view channels.

**Calibrate**

If a waveform, WaveMark or RealWave channel has sections with known amplitudes, offsets, slopes or areas, you can calibrate it. Spike2 supports a wide variety of calibration methods for these channels. You can calibrate a single channel, or all selected channels (as long as they have the same calibration values). If you make a mistake, you can **Undo** the calibration changes. You can also calibrate from a script.



To calibrate your data, click on a time window and open the calibrate dialog. If there are any suitable channels selected, the channel list in the dialog shows the selected channels, otherwise you must choose a channel for the drop down list. Once the dialog is open you can select channels in the time window; these are added to the selected list in the dialog.

Now choose a calibration method from the list. The dialog contents change depending on the selected method. All methods require you to select data areas with known values; this is most easily done with cursors. Click **Set cursors** to position the appropriate cursors in the time window and dialog. You can also type in the times and use expressions like $(\text{Cursor}(1)+\text{Cursor}(2))/2$. If the method requires two time ranges they must not overlap.

When you change the method or the channel or click the **Update** button, Spike2 collects the current calibration values from the (first) channel in the **Channels** field.

Once you have selected your data you must type in the calibration values, and also set the units of the data. A text box displays Use OK or Apply to calibrate or an explanation of any problem that prevents calibration. Click Apply to calibrate and leave the dialog open or OK to calibrate and close the dialog.

If you check the Set sampling configuration box, any calibration changes are passed through to the same channel number in the sampling configuration. If you calibrate with a data file that you are sampling, or that you have just sampled and not yet saved, this box is checked automatically when the dialog opens.

What calibration does Waveforms stored on disk as 16-bit integers (range -32768 to 32767) are scaled into user units by a scale and offset: $\text{User units} = 16\text{-bit integer} * \text{scale} / 6553.6 + \text{offset}$

The factor 6553.6 is present so that if the input data range spans -5 to +5 Volts (as is usually the case with a CED 1401 interface), the relationship is:

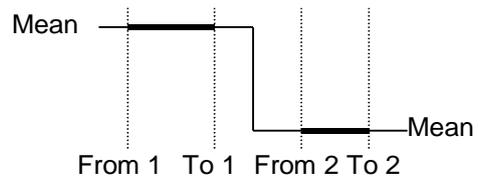
$$\text{User units} = \text{input in Volts} * \text{scale} + \text{offset}$$

The scale and offset are the same values as you set in the sampling configuration dialog or in the channel information dialog in a time window. Calibration changes the scale and offset so that the displayed data matches your user units.

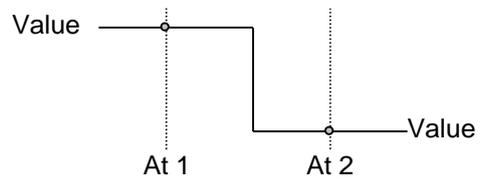
RealWave data is stored as floating point numbers on disk. Calibration rewrites this data and so may take noticeable time when working on very large data files. You cannot calibrate a RealWave data if any channel process is attached. You cannot calibrate any channel with an attached process that has changed the channel scale or offset value.

Calibration methods The Method field of the calibration dialog sets how the calibration is performed:

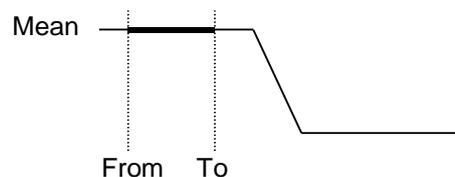
Mean level of two time ranges You define two time ranges in your data (From 1 to To 1 and From 2 to To 2) and supply the mean levels of the two ranges (Mean 1 and Mean 2). The data in the time ranges must have significantly different levels and the two mean values you supply must not be the same.



Values at two times You define two times (At 1 and At 2) and the two values (Value 1 and Value 2) that correspond to the two data values. The two values must be different and the data at the two times must also be different.

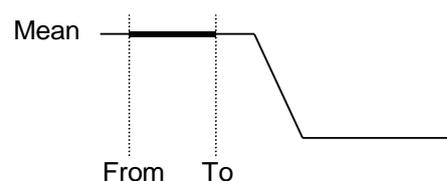


Set offset from mean of time range If your data is calibrated, but suffers from baseline drift, you can use this method to redefine your base line. Set a time range (From and To) and the mean value of the data in the range (Mean). The data scaling is not changed.



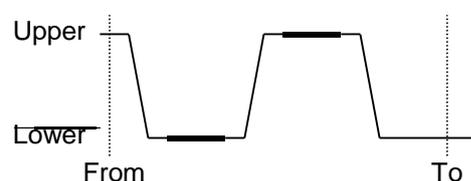
Set scale from mean of time range

This method is for data with a fixed offset (usually 0) and variable gain. Set a time range (From and To) and the mean value of the data in the range (Mean). The data zero level is preserved. The mean value you set cannot be 0 and the mean level of the data before calibration cannot be 0.



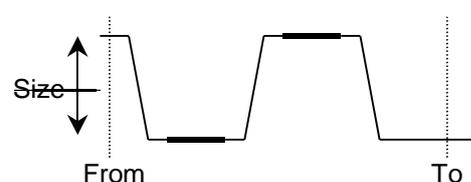
Square wave, upper and lower level

This method will calibrate a square wave with known Upper and Lower levels. The time range (From and To) must contain at least three transitions between levels. In this case Upper means the larger value before calibration. To calculate the mean value of these levels, Spike2 detects the transition points between high and low values. For each high or low section, the first and last 25% of the data is ignored. The upper and lower levels must not be the same.



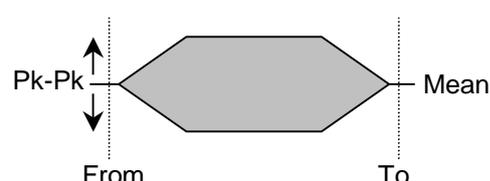
Square wave, amplitude (Size) only

This method detects a square wave in the same way as the *Square wave upper and lower level* method. In this case, you supply the Size (difference between the upper and lower levels) of the waveform. The zero level is preserved.



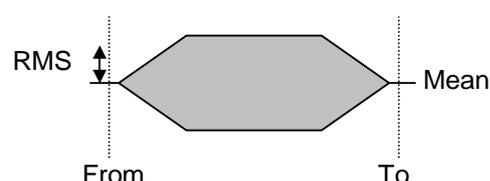
Peak to peak amplitude and mean

This method calibrates based on the peak to peak amplitude (Pk-Pk) and Mean value of the data in the time range (From and To). This could be used for a sinusoidal waveform of known amplitude. The Peak to Peak value must be non-zero.



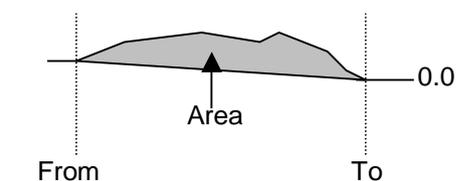
RMS amplitude about mean

This method calibrates based on the RMS (Root Mean Square) amplitude of the data in the time range (From and To) and the Mean level. The RMS value must be non-zero.



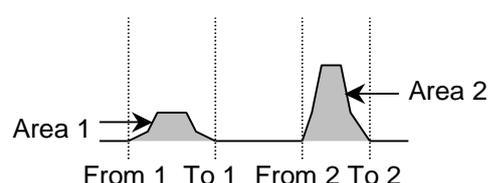
Area under curve, assume zero at end

You could use this method to calibrate a rate based on a known area (for example flow rate based on volume). The rate is assumed to be zero at each end of the time range (From and To) to allow for a drifting base-line. The area you set is in the units of the rate units times seconds. The rate value at the To time is set to zero (if the From value is different the base line is assumed to run linearly from From to To).



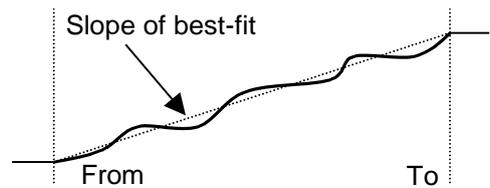
Areas under curve, two time ranges

Use this method when you know the area under the data for two sections (From 1 to To 1 and From 2 to To 2) of your data. The mean level of the two sections must be different. The area you set is in the units of the channel units times seconds.



Set scale from slope (no offset change)

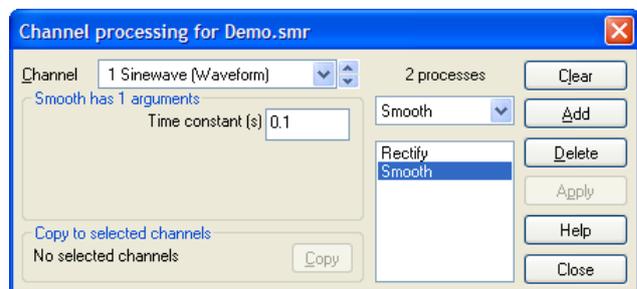
Use this method when you know the slope of a section of the data (for example calibrate position based on a known velocity, or velocity based on a known acceleration). The calibration is set so that the slope of the best-fit line to the data in the range matches the slope value you set. This method does not change the offset, so you may need to combine this method with the *Set offset from mean of time range* method for a full calibration. The units of the slope you set are the channel units per second.

**Channel process**

A channel process is an operation, for example rectification, applied dynamically to waveform or RealWave data channels. The original data is not changed, but all users of the channel see data modified by the process. Multiple processes can be applied, for example DC removal then rectification then smoothing. Every time you use the data, the processing is applied, so using a processed channel is slower than using raw data.

If the wave is already part of an analysis operation, such as a waveform average, adding a process and continuing the analysis may generate incorrect results. This is particularly the case with processes that change the sample rate or interpretation of the data.

To add a process, select the analysis menu **Channel Process** command and select a **Channel**. The list to the left of the **Apply** button shows processes to apply in order. Values in the selected process that you can change are displayed in the box on the left of the dialog. The



Clear button removes all processes from the channel. **Add** appends a new process set by the drop down list to the left of the **Add** button to the end of the list. **Delete** removes the selected process and **Apply** copies changed argument values to the process.

Copy to selected channels

You can copy all the processes set for the current channel to other channels. Select the target channels by clicking their channel numbers (usually on the left edge of the time view) and click the **Copy** button. If the current channel has no processes, this will clear all processes from the selected channels.

Restrictions caused by processing

Some processes change the channel scale and offset (these are the values that translate between a 16-bit integer representation of a waveform and user units). Such changes do not affect the data on disk and are removed when the process is removed. The **Calibrate** dialog and the **Channel Information** dialog will not allow you to change the channel calibration if an attached process has changed the channel scale or offset. You are not allowed to calibrate a RealWave channel that has any attached channel process.

You can add processes of the following types:

Rectify

This replaces all negative input values with positive values of the same magnitude. The result of this operation may exceed the 16-bit range of a waveform channel if the channel offset is negative, in which case the output will be limited to the available 16-bit range. There are no additional arguments required to define this process.

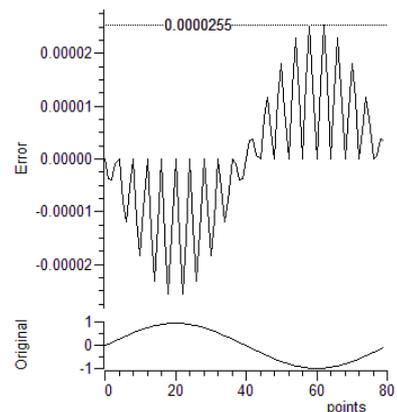
- Smooth** This process has one argument, a time period in seconds, p . The output at time t is the average value of the input data points from time $t-p$ to $t+p$ seconds.
- DC Remove** This process has one argument, a time period in seconds, p . The output at time t is the input value at time t minus the average value of the input data points from time $t-p$ to $t+p$. This process does not affect the channel scale, but the channel offset is set to zero.
- Slope** This process has one argument, a time period in seconds, p . The slope at time t is calculated using an equal weighting of the points from time $t-p$ to $t+p$. If you apply this process to a channel, the channel scale, offset and units change. If the current channels units are no more than 3 characters long, "/s" is added to them, so units of "V" become units of "V/s". If there is not sufficient space, the final character of the units becomes "!" to indicate that the units are no longer correct. The offset becomes 0, and the scale changes to generate the correct units.
- Time shift** This process has one argument, the time to shift the wave. A positive time shifts the wave into the future (to the right); a negative time shifts the wave into the past. If you want to keep the waveform unshifted, but change the positions where the data was sampled, use the Interpolate or Match channel processes.
- Down sample** This process changes the sample rate of the wave by taking one point in n . There is one argument, prompted by Use one point in, which is the down sample ratio. You might want to use this command after filtering or smoothing a waveform.
- Interpolate** You can change the sample rate of a channel and set the time of the first data point with this process. Interpolation is by cubic splining the original data. No data is generated outside the time range of the original data points. Interpolation is not too slow, but if you increase the sampling rate it will take longer to draw and process data.

The first argument, **Sample interval**, is the time in seconds between output data points. You can type in expressions here so, to set 123 Hz, type $1/123$. The actual interval is set as close to the requested one as possible. When you create the process this is set to the current sample interval of the channel.

The second argument, **Align to**, aligns the output data to a time. It must be positive. The process places data points at this time plus and minus multiples of the sample interval. You can use this to convert multi-channel data sampled by a 1401 into simultaneously sampled data by giving all the channels the same alignment and sample rate.

Cubic spline interpolation assumes that the input waveform and its first and second differentials are continuous. If the input data was suitably filtered this will be not too far from the truth. Not all data is suitable for cubic splining; splining across step changes generates ringing effects that were not present in the original signal.

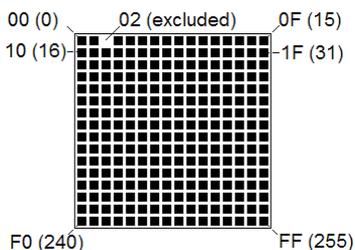
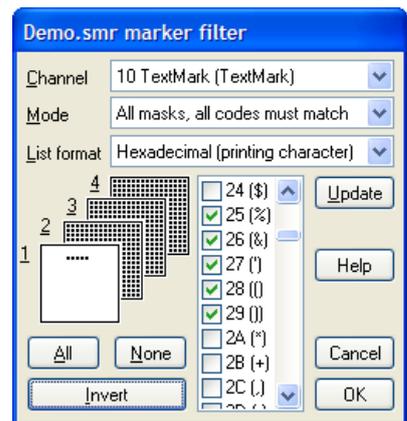
Cubic spline interpolation is better than linear interpolation for continuous data, but it is not perfect. The graph shows the error between a sine wave sampled with 20 points per cycle and splined to 80 points per cycle and a calculated 80 points per cycle sine wave. The maximum error is small, in this case 0.0000255 of the original signal. However, the maximum error increases rapidly the fewer points there are per cycle of the original data. With 5 points per cycle, the maximum interpolation error for a sinusoid is almost 1 per cent of the original signal.



- Match channel** This is the same as *Interpolate*, except that the sample interval and alignment are copied from a nominated channel. The initial channel is set to the current channel, so adding this process should have no visible effect (apart from causing a redraw).
- RMS amplitude** This process has one argument, a time period in seconds, p . The output at time t is the RMS value of the input data points from time $t-p$ to $t+p$ seconds. For waveform data, the output may be limited by the 16-bit nature of the data if the channel offset is non-zero.
- Median filter** This process has one argument, a time period in seconds, p . The output at time t is the median value of the input data points from time $t-p$ to $t+p$ seconds. The median is the middle point after the data has been sorted into order. This can be useful if your data has occasional points with large errors. This filter is slow if p spans a large number of data points; set the time period to the smallest value that removes the outlier points.
- Fill gaps** Waveform and RealWave channels can have gaps, where no data exists. This process guarantees that the channel has no gaps by filling in gaps greater than a specified time with a fixed level and by linearly interpolating across smaller gaps. There are two arguments: the maximum gap in a waveform to interpolate across and the level to use to fill gaps wider than the maximum gap. At the start and end, gaps less than the specified maximum are filled by duplicating the first or last data point.
- Skip NaN** A NaN (Not a Number) is a floating point value that is either undefined in some way, or is infinite (the result of dividing by zero). This can occur when RealWave data is imported into Spike2. This process removes such numbers from the channel, leaving gaps. Applying this process to a Waveform channel is a waste of processing time. The process has no arguments.

Marker Filter You can filter any channel that holds marker, WaveMark, TextMark or RealMark data. Each of these channels has a *marker filter* that selects the data items to display and use in calculations. Each marker data item has four marker codes that are matched against the marker filter.

Masks The dialog shows the marker filter as four masks numbered 1 to 4. Each mask has 256 elements in a 16x16 grid, one for each possible code value. The contents of the front-most mask are displayed in the scrolling list in the centre of the dialog. Click on a mask element to bring the mask to the front and scroll the list to the element.



The top row of each mask represents code values 0 to 15 (hexadecimal codes 00 to 0F), the second row 16 to 31 (10 to 1F) and so on down to the bottom row, which represents values 240 to 255 (F0 to FF). If a code value is included in the filter, the corresponding element is black. When values are excluded, the element is white. The mask gives a quick indication of the state of the filter. There are three buttons that act on the entire mask:

- All** This includes all code values in the filter (checks all the list boxes)
- None** This excludes all code values from the filter (clears all the list check boxes)
- Invert** This excludes included values and includes excluded values

Channel This field is the standard channel selector. You can select only marker-based channels.

List format The list shows each the code value as two hexadecimal or one to three decimal digits, the single character equivalent or as a combination as set by the *List format* field. Characters are appropriate for keyboard markers, numerical digits are used for the non-printing codes or can be forced for all codes. If you type a character, the list scrolls to the entry that starts with the typed character. To include marker codes, check the boxes. There are two modes in which to use the marker filter:

Mode 0: All masks, all codes must match For a data item with marker codes *a*, *b*, *c* and *d* to be included, mask 1 must have code *a* checked, mask 2 must have code *b* checked, mask 3 must have code *c* checked and mask 4 must have code *d* checked. Most users of this mode set mask layers 2, 3 and 4 to All and use the first layer to select data values. You can think of this as the *and* mode; to accept data marker code 1 must be in the layer 1 *and* marker code 2 must be in the layer 2 *and* marker code 3 must be in the layer 3 *and* marker code 4 must be in layer 4.

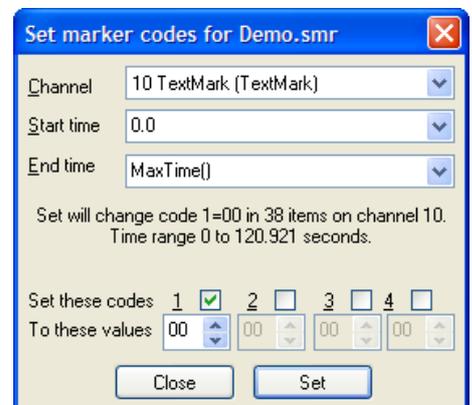
Mode 1: One mask, any code can match Only mask 1 is used, the rest are greyed out. For a data item with marker codes *a*, *b*, *c* and *d* to be included, mask 1 must have one or more of the codes *a*, *b*, *c* and *d* checked. You can think of this as the *or* mode; to accept data marker code 1 *or* marker code 2 *or* marker code 3 *or* marker code 4 must be in the layer. There is once exception; for marker codes 2 to 4, the code 00 is ignored. To accept code 00, it must be the first marker code.

Mode 1 is often used when sorting spike shapes (WaveMark data) and you discover a WaveMark that is the result of a collision between two spikes. You can set the first marker code to the code for the first spike and the second to the code for the second (leaving the third and fourth codes as 00), then the spike will appear on screen and in analyses when either of the codes are included in the mask.

Cancel, Update, Help, OK The buttons on the right are:
Update Update the channel display to correspond with the new filter
Help Display this information
Cancel Close dialog, cancelling changes since last update
OK Close dialog, accept new filter

Set Marker Codes

You can open this dialog from the Analysis menu, by right clicking on a marker channel and from the Edit WaveMark dialog. From here you set selected marker codes for marker, WaveMark, TextMark and RealMark channels. Each data item has four marker codes that can be used to filter the data in the channel. Usually only the first marker code is used, see the Marker Filter command for more information. The Channel field sets the channel to process and the Start time and End time fields identify the markers to be set. If a marker filter is set for the channel, only data items that are in the filter are changed by this command. The four check boxes 1 to 4 select the codes to change. You set the value for each code as either two hexadecimal digits or one printing character.



Example To change all the marker items in a channel with a first marker code of 02 to have a first marker code of 1A do the following:

1. Open the marker filter dialog and set the channel to display only marker code 02.

2. Open this dialog, select the channel and set the time range as 0.0 to Maxtime().
3. Check the 1 box to set the first marker code and edit the text under the check box to 1A. Make sure the other check boxes are clear.
4. Read the text in the centre of the dialog as a check that this is the action you wish to take. This action cannot be undone.
5. Click the Set button to change the marker codes.

All the code 02 items will disappear as they have been coded as 1A, so you will need to change the marker filter settings if you want to see the result.

New WaveMark This creates WaveMark data channels from waveform data or existing WaveMark data (see the *Spike shapes* chapter for details). You can also activate this option by right-clicking on a waveform or WaveMark channel and selecting the New WaveMark command from the context menu.

New n-trode Use this to create a new WaveMark channel with 2 or 4 traces. To enable this option, select 2 or 4 waveform channels with the same sample rate. This option opens the New WaveMark dialog with the selected channels as the source (see the *Spike shapes* chapter for details). The channels are used in the order that they are selected. The first selected channel sets the sample rate that the selected channels must match.

Edit WaveMark This option reclassifies WaveMark data both manually and automatically (see the *Spike shapes* chapter for details). You can also activate this option by right-clicking on a WaveMark channel and selecting the Edit WaveMark command from the context menu.

Digital filters This displays a pop-up menu in which you can choose the FIR Digital filtering dialog or the IIR Digital filtering dialog. From these dialogs you can create digital filters and apply them to waveform or RealWave channels (see the *Digital filtering* chapter for details). You can also open these dialogs by right clicking on a suitable channel and selecting a digital filter option from the context menu.

Window menu

The **Window** menu controls the data and text windows that belong to the Spike2 application. It has commands to duplicate a time window, commands to hide and show windows, a command to close all windows and commands to arrange the windows within the application window. The remaining space at the bottom of the menu holds a list of all the windows that belong to the Spike2 application. If you select one of the windows in the list, the window is brought to the front and made the current window.

Duplicate window This command is only available within a time window and creates a duplicate window with all the attributes (list of displayed channels, event display modes, colours, cursors and size) of the original window. Once you have created the new window, it is independent of the original. However, the data channels within it are the same data channels as in the original window, so any changes made to the data in one window will cause all duplicated windows to update. Duplicating a window allows you to have different views of the same data file with different scales and drawing modes and different sets of channels.

You can close all windows associated with a data document using the control key plus close (see the **Close** command in the *File menu* chapter). This will remember the position and state of all windows associated with the document.

Hide This command makes a window invisible. This is often used with script windows and sometimes is used to hide data windows during sampling when only the result views are required. Closing the Log window is equivalent to hiding it as the Log window always exists.

Show This option lists all hidden windows in a pop-up menu. Select a window from the list to make it visible.

Tile Horizontally This command arranges all the visible screen windows so that no window is overlapped by any other. Iconised windows are arranged along the bottom edge of the window. The command attempts to arrange window so that they are wider than they are high. Tiling takes into account the space used for title bars of iconised windows, so to use the full application window area you should hide any iconised windows first. This command may have the same result as **Tile Vertically**, depending on the number of windows.

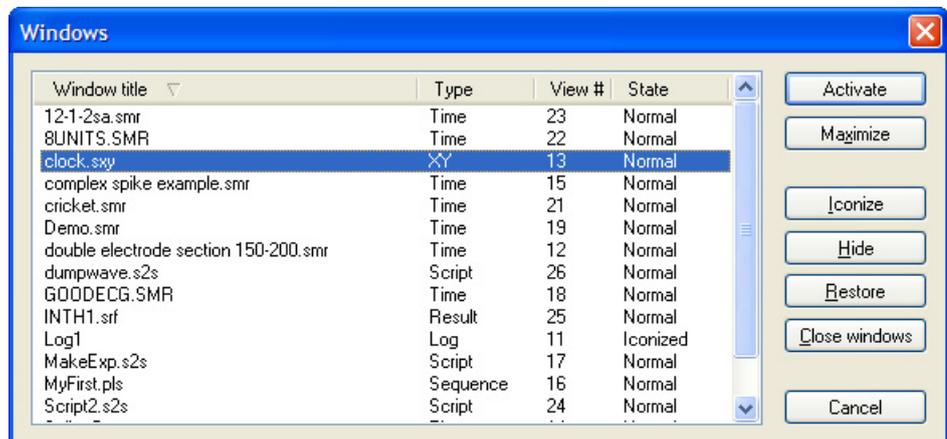
Tile Vertically This command arranges all the visible screen windows so that no window is overlapped by any other. Iconised windows are arranged along the bottom edge of the window. The command attempts to arrange windows so that they are taller than they are wide. Tiling takes into account the space used for title bars of iconised windows, so to use the full application window area you should hide any iconised windows first.

Cascade All windows are set to a standard size and are overlaid with their title bars visible. Any iconised windows are left in the iconised state, and they are arranged along the bottom edge of the window, as for the **Arrange Icons** option.

Arrange Icons You can use this command to tidy up the windows that you have iconised in Spike2. The icons are lined up along the bottom edge of the application window.

Close All This command closes all windows in the Spike2 application. You are asked if you want to save the contents of any text windows that have changed or any newly sampled data window. You can avoid being asked if you want to save modified result and XY windows with an option in the Edit menu Preferences. The positions of data windows and attached result view windows are saved.

Windows This dialog lists all the document-related windows that are open and lets you apply common window operations to one or more of the windows. You can sort the list based on the window title, type, view number (as seen by the script language) and window state by clicking the title bar at the top of the list.



Cursor menu

A cursor is a vertical or horizontal line drawn in a time, result or XY view, to mark or obtain a position. The **CURSOR** menu creates and destroys cursors, changes the display to make them visible, changes their labelling mode and obtains the values of channels where they cross the cursors and between the cursors. Up to 10 vertical and 9 horizontal cursors can be active in each window. Cursors can be dragged over and past each other. If you duplicate a time view the cursors are also duplicated, but are independent of the originals.

In addition to using the Cursor menu commands, if you right-click on a vertical or horizontal cursor there are additional cursor commands available from the context menu. For example you can copy the position of cursor, or the difference of two horizontal cursor positions to the clipboard.

New Cursor



This command and the cursor button at the bottom left of time and result windows adds a vertical cursor with the current label style and lowest available cursor number (1 to 9) to the window. The keyboard short cut `Ctrl+n` where `n` is a cursor number (including 0 in a time window) creates cursor `n` if it doesn't exist and moves it to the centre of the window. Each view has a default label mode for new cursors that is set by the **Label mode** command.

Delete

The delete command activates a pop-up menu from which you can delete one or all vertical cursors. As an aid to identification cursors are listed with their number and position. Deleting a cursor removes it from the window; other cursors are not affected.

Fetch

This opens a pop-up menu in which you select a vertical cursor to place in the visible x axis. The keyboard short cut `Ctrl+n` where `n` is a cursor number will fetch cursor `n` if it exists or will create it if it doesn't exist.

Move To

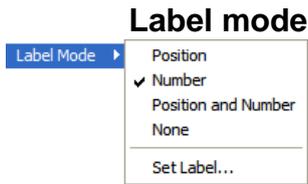
This command opens a pop-up menu from which you can select a vertical cursor to move to. The cursors are listed with their number and position as an aid to identification. The window is scrolled to display the nominated cursor in the centre of the screen, or as close to the centre as possible. This command does not change the x axis scaling. You can also use the keyboard short cut `Ctrl+Shift+n` where `n` is a cursor number.

Position Cursor

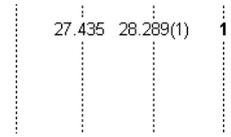
This command opens a pop-up menu to select a cursor and then opens a dialog in which you can type or select a cursor position. You can also activate this dialog by right clicking on a vertical cursor and selecting **Set Position** from the cursor context menu.

Display All

This command has no effect if there are no active cursors. If there is a single cursor, the command behaves as though you had used the **MOVE TO** command and selected it. When there are multiple cursors, the window is scrolled and scaled such that the earliest cursor is at the left-hand edge of the window and the latest is at the right-hand edge.



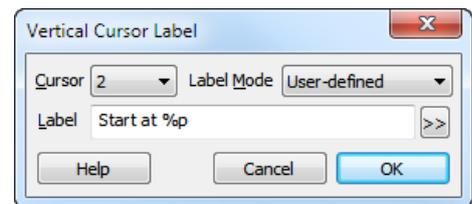
Each cursor has an optional label used to identify it. You can drag the cursor labels up and down the cursor with the mouse to suit the data. There are five cursor label modes: **None**, **Position**, **Number**, **Position and Number**, and **User-defined**. You select the most appropriate for your application using the pop-up menu or by right clicking on a cursor and choosing to set the cursor label from the context menu. To avoid confusion between the cursor number and the position, the number is displayed in bold type when it appears alone and bracketed with the position.



Each cursor stores its own mode and label, and the view has a mode that is applied to new cursors. The first four items in the menu set the view mode and the mode of all cursors. You can also set a user-defined label for cursors (but not for the view) with the **Set Label** command.

Set Label

You can open the cursor label dialog from the **Label Mode** and **Horizontal Label Mode** cursor menu commands or by right-clicking on a cursor and using the **Set Label** command from the cursor pop-up menu.



From this dialog you can set the cursor label mode for one or all cursors. The **Cursor** field can be set to the number of any cursor in the view, or **All**. If you choose **All**, any change applies to all cursors and sets the view mode except in **User-defined** mode, where the view mode does not change.

The **Label Mode** field lets you choose one of **None**, **Position**, **Number**, **Position and Number**, and **User-defined** as the cursor mode. If you select **User-defined**, the **Label** field appears together with the **>>** button and you can set a label of your choosing.

User-defined labels

User-defined labels display the text you type, except that the text sequences $\%p$, $\%u$, $\%n$ and $\%v(n)$ are replaced with the cursor position, cursor units, cursor number, or the value of channel n where it is crossed by the cursor. You can also stipulate the width (w) and the number of decimal places (D) used for the position and value by using $\%w.Dp$ and $\%w.Dv(n)$. For example: $\%n$ at $\%6.4p$, $\%v(2)$ might display: **1** at 2.2346, 87.128756 if cursor 1 was at 2.2346 x axis units and channel 2 had the value 87.128756 at this point. The $\%v(n)$ option is not allowed for horizontal cursors or for vertical cursors in an XY view. The value returned by $\%v(n)$ is the same value as displayed in the cursor values dialog for that cursor and channel.

The **>>** button pops up the list of replacements, and if you choose one, it replaces any selection in the **Label** field. If you choose the $\%v(n)$ option, you are prompted to select a channel to measure.

We allow you to type in quite long labels. However, when you close a data file, only the first 19 characters of a user-defined label are saved and long labels look messy, so it is usually a good idea to keep labels short.

Renumber

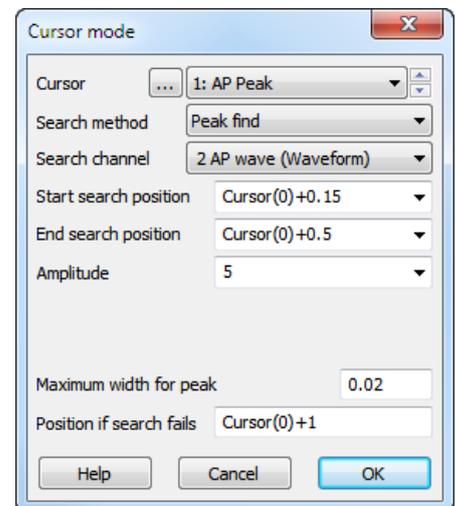
When created, cursors take the lowest available cursor number. You can also drag cursors over each other. This command renumbers the vertical cursors, with cursor 1 on the left.

Active cursors In a time view, vertical cursors can be *active* or *static*. An active cursor can seek to a position based on the position of other cursors and data in a channel. Active cursors can automate data analysis, leading to new data channels, XY trend plots or tabulated output.

Cursor 0 Vertical cursor 0 is special. It always exists and cannot be deleted, but it can be hidden. It is the iterator for XY trend plots; a movement of cursor 0 causes all other active cursors to recalculate their positions. This calculation is done in order of rising cursor number. To hide cursor 0, right click on it and select **Hide cursor 0** in the context menu.

Valid and invalid cursors Active cursor positions are either *valid* or *invalid*; invalid cursors have an exclamation mark at the end of the label. A position is invalid if a search fails and the **Position if search fails** field is empty or does not contain a valid expression. Expressions that use invalid cursor positions are also invalid. The XY Trend plot rejects points with invalid measurements. Cursor positions are validated by operations that move them to a specific place such as dragging. If a search results in an invalid position, the cursor is not moved.

Active mode This command opens the Cursor mode dialog. The **Search method** field and the selected cursor determine the dialog contents. The ... button opens the **Cursor Label** dialog to set a user-defined label. An active cursor has an associated **Search channel** and start and end positions that define the data to search to locate the new cursor position.



Cursor 0 does not have start and end positions. However, it does have a **Minimum step**; searches start at this distance from the cursor 0 position and continue to the file end for a forward search, or to the file start file for a backward search. Cursor 0 has a restricted range of search methods: **Peak find**, **Trough find**, **Rising threshold**, **Falling threshold**, **Within thresholds**, **Outside thresholds**, **Peak slope**, **Trough slope**, **Turning point**, **Data points** and **Expression**.

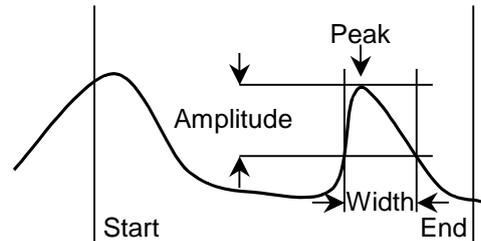
The search positions can be a fixed time, but usually they are expressions that involve the positions of other active cursors. Active cursor positions are evaluated in sequence from cursor 0 to 9. For cursor n , an expression that refers to an active cursor less than n uses the new position. An expression that refers to a cursor greater than or equal to n uses the old position.

If the **End position for search** is less than the **Start position for search**, searches are backwards. If a search is backwards, read *previous* for *next* and *last* for *first* in the descriptions of the cursor modes. Where search modes depend on data values, for example maximum and minimum, and the data channel does not have a y axis, then the value is taken as the interval between data items.

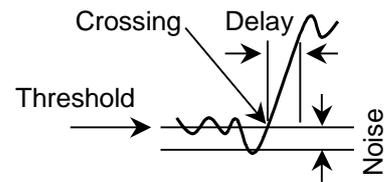
Several modes use a slope. These modes have the value **Width for slope measurement**, in seconds and use the data points from $Width/2$ before the current position to $Width/2$ after to calculate the slope unless there are more than 200 points, in which case 100 points before and after are used. The contribution of each point to the slope is proportional to the distance of the point from the current position. Because the slope at any point requires data around it, the slope within $Width/2$ of the ends of the data and any gaps is not to be relied on as the missing data is assumed to be zero.

Static A new cursor starts out in **Static** mode. It stays where you put it and is not changed by a change in the position of a lower numbered cursor. The cursor position is always valid.

Peak find, Trough find The **Amplitude** field defines how much the data must rise before a peak and fall after it (or fall before a trough and rise after it), to detect a perk/trough. The **Maximum width for peak field** rejects peaks that are too broad (set it 0 for no width restriction). For waveform channels, the peak position is located by fitting a parabola through the highest point and the points on either side. In the diagram, which shows a peak search, the first peak is not detected because the data did not rise by **Amplitude** within the time range. The cursor position is valid if a peak is detected, invalid if not.



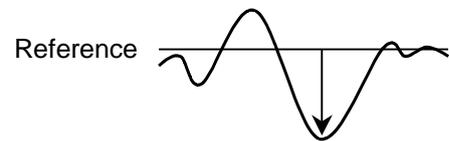
Rising threshold, Falling threshold, Threshold The data must cross **Threshold** from a level that is more than **Noise rejection** (hysteresis) away from it and stay crossed for a time of at least **Delay** after level crossing. For the **Rising threshold** mode, the data must increase through the threshold, for **Falling threshold** mode it must fall through the threshold. In **Threshold** mode the crossing can be in either direction. The picture shows a rising threshold. For waveform channels, the crossing point is found by linear interpolation of the data points on either side of the threshold crossing. The cursor position is invalid if a crossing does not occur.



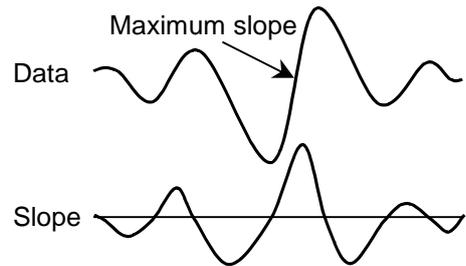
Outside thresholds, Within thresholds These two modes are similar to the rising and falling thresholds modes except they search for data the lies outside two threshold levels or within two threshold levels. In addition to the **Threshold** level and **Noise rejection** (hysteresis) fields, there is a **Threshold level 2**, which sets the second threshold level. The data must cross a threshold from a level that is more than **Noise rejection** away from it.

Maximum value, Minimum value The result is the position of the maximum or minimum value. The result is invalid if there is no data in the search range.

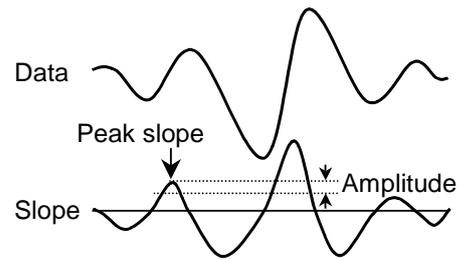
Maximum excursion There is an extra field in this mode for the **Reference** level. The cursor is positioned at the point that is the maximum distance in the y direction away from the reference level. The result is invalid if there is no data in the search range.



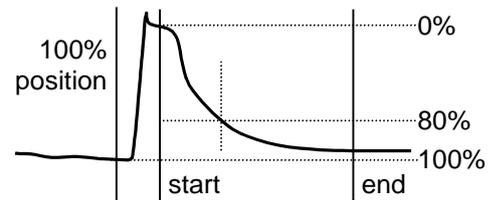
Steepest rising, Steepest falling, Steepest slope (+/-) These modes are for waveform channels only. They have the extra field **Width** for slope measurement that sets length of data used to evaluate the slope at each data point. The result is the position of the maximum, minimum or maximum absolute value of the slope. The result is invalid if there is no data or not enough data to calculate a slope or if the channel is not waveform or WaveMark.



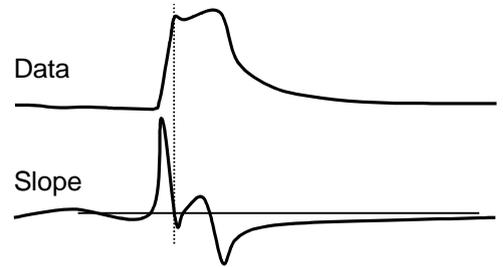
Slope peak, Slope trough These analysis modes find the first peak or trough in the slope within the search range that meets the **Amplitude** specification. The **Width** for slope measurement field sets the length of data used to evaluate the slope at each data point. The **Amplitude** field sets how much the slope must rise before a peak and fall after it (or fall before a trough and rise after it), for it to be accepted as a peak. The **Amplitude** units are y axis units per second.



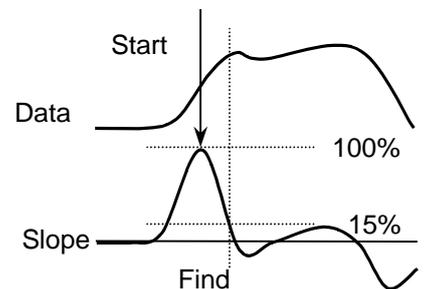
Repolarisation % This mode finds the point at which a waveform returns a given percentage of the distance to a baseline inside the search range. The start of the search range defines the position of 0% repolarisation. The 100% position, which can be anywhere, and **Width** fields identify the 100% level. The **Repolarise %** field (drawn at 80% in the picture) sets a threshold level in percent relative to the 0% and 100% levels. The position is the first point in the search range that crosses the threshold. The result is invalid if the threshold is not crossed.



Turning point This finds the first point in the search range where the waveform slope changes sign. Put another way, it finds a localised peak or trough. The **Width** for slope measurement field sets the data range to calculate the slope. The picture shows how this can find the top of a sharp rise where **Maximum** would get the wrong place. To use this, set a cursor on the peak slope and start the search from that point. The result is invalid if no point is found.



Slope% Use this to find the start and end of a fast up or down stroke in a waveform. The **Width** for slope measurement field sets the time width used to calculate the slope. The **Slope%** field sets the percentage of the slope at the start of the search area to find. Set a cursor on the maximum or peak slope, then use that as the start point and search for the required percentage. A value of 15% usually works reasonably well. The result is valid if the slope value is located.



Data points The **Point count** field sets the number of data points to move by in the search range on the referenced channel. The result is the time of this data point. The cursor position is invalid if there is no data item in the search range.

Expression The new position is obtained by evaluating the **Position** field, which normally holds an expression based on cursor positions, for example `Cursor(0)+1` or `(C0+C1)/2`.

**Search Right,
Search Left**

If cursor 0 is active, these two command cause the cursor to search for the next or previous position that satisfies the active mode. If the cursor mode is **Expression**, the cursor goes the same way for both commands.

Display y values

This command opens a new window containing the values at any cursors in the current time or result window. Cells for cursors that are absent, or for which there is no data, are blank. The new window is a top level window and will never go behind another time or result view.

Cursors	Cursor 0	Cursor 1	Cursor 2	Cursor 3
Time (s)	26.428747	30.807193	35.18564	39.564087
31 DupKey	27.03616	31.92832	37.04064	39.65696
1 Sinewave	2.5869173	-0.40142127	0.22451212	2.960858
5 Response	19.8036	18.549735	16.170691	16.120529
2 Stimulus	26.44303	30.82886	35.20149	39.61393
<input type="checkbox"/> Time Zero	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
<input type="checkbox"/> Y Zero	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

In a result window, the value displayed is the value to be found in the bin to the right of the cursor. If you position the cursor at the far right (where there is no bin to the right), a blank is displayed as the value.

In a time window, the values displayed depend on the channel display mode. If a channel has a y axis the displayed value will be in y axis units. If the data is displayed as a continuous line or as a histogram, the displayed value is where the line crosses the cursor. If the data is displayed as dots, the value is the nearest dot for waveforms and the previous dot for instantaneous frequency. Positions more than one waveform sample interval away from any waveform point have no value.

If the channel has no y axis, in State drawing mode, the value is the state. Otherwise the channel holds an event or marker and the displayed value is the time of the next event or marker at or to the right of the cursor. If there is no data, the field is blank.

The **Time zero** check box enables relative cursor times. If checked, the cursor marked with the radio button is taken as the reference time, and the remaining cursor times are given relative to it. The reference cursor displays the absolute time, not 0. In the example above, cursor 1 has been set as the reference.

The **Y zero** check box enables relative cursor values. The radio buttons to the right of the check box select the reference cursor. The remaining channels display the difference between the values at the cursor and the values at the reference. The values for the reference cursor are not changed.

Selecting, copying and printing data

You can select areas of this window and copy them to the clipboard by clicking on them with the mouse. Hold down the **Shift** key for extended selections. You can select entire rows and columns by clicking in the cursor and channel title fields. Hold down the **Control** key to select non-contiguous rows and columns. Use **Ctrl+C** to copy selected data to the clipboard. Use **Ctrl+L** to copy selected data to the log view. Right click in the cursor window to display a context menu with command for Copy, Log, Font and Print.

Cursor regions



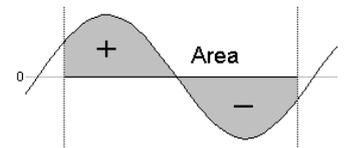
This command opens a cursor region window for the current time or result view. This window calculates values for data regions between the cursors. One column can be designated the Zero region by checking the box and selecting the column with a radio button. The value in this column is then subtracted from the values in the other columns (except in Area(scaled) mode which scales the zero region value to allow for column widths). The field at the bottom left indicates how to calculate the values; click on it for a full list. Cells for which there are no cursors or data are blank. You can interrupt long calculations in time windows with the **Ctrl+Break** key combination.

Cursors	0 - 1	1 - 2	2 - 3
Time (s)	5.0247856	4.3784467	4.3784467
1/Time (Hz)	0.19901347	0.2283915	0.2283915
31 DupKey	0.39802738	0.45678265	0.22839133
1 Sinewave	1.3717608	1.3835991	1.3982419
3 Response	17.862161	21.821886	14.582918
2 Stimulus	13.532931	18.728089	18.271306
<input type="checkbox"/> Zero Region	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Mean	<input type="text"/>		

Time window waveform data and result window data

In a time window, waveform channels include WaveMark channels drawn as waveforms, WaveMark or Cubic Spline and waveform and RealWave channels drawn as dots, Cubic spline or Sonogram. The region set by a pair of cursors is the data starting at the first cursor up to, but not including, the data at the second cursor. In a result window, the region set by a pair of cursors starts at the bin containing the left cursor and ends in the bin to the left of the bin containing the right cursor. The values for each mode are:

Area The area between the data points and the y axis zero. Area is positive for curve sections above zero and negative for sections below zero. Use Modulus if you want areas below the y axis to be treated as positive.



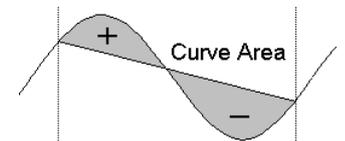
Mean The sum of all the data points between the cursors divided by the number of data points between the cursors. There is also a Mean Absolute value.

Slope The slope of the least squares best fit line to the data between the cursors.

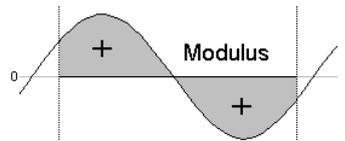
Sum The sum of the data values between the cursors. If there are no samples between the cursors the field is blank.

Area (scaled) The same as Area, but if a zero region is specified, the amount subtracted from the other regions is scaled by the relative width of the regions.

Curve area Each data point makes a contribution to the area of its amplitude above a line joining the endpoints multiplied by the x axis distance between the data points. The picture makes this clearer.



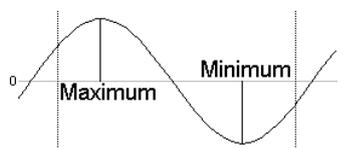
Modulus Each data point makes a contribution to the area of its absolute amplitude value multiplied by the x axis difference between data points. This is equivalent to rectifying the data, then measuring the area. If a zero region is specified, the amount subtracted from the other regions is scaled by the relative width of the regions.

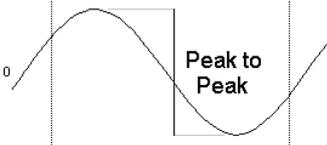
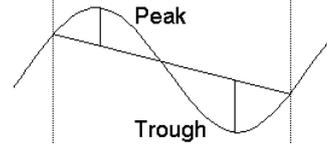
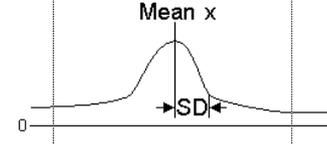


Maximum The value shown is the maximum value found between the cursors.

Minimum The value shown is the minimum value found between the cursors.

Abs Max The value shown is the maximum absolute value found between the cursors. If the maximum value was +1, and the minimum value was -1.5, then this mode would display 1.5.



Peak to Peak	The value shown is the difference between maximum and minimum values found between the cursors. The peak to peak value is always positive.	
SD	The standard deviation from the mean of the values between the cursors. If there are no values between the cursors the field is blank. If there are n data points, and the sum of the squares of the differences between the points and the mean value is $DiffSquare$, the result is calculated as $Sqrt(DiffSquare / (n-1))$.	
RMS	The value shown is the RMS (Root Mean Square) of the values between the cursors. If there are no values between the cursors the field is blank.	
Peak	The value shown is the maximum value found between the cursors measured relative to a baseline formed by joining the two points where the cursors cross the data. This is always greater than or equal to 0.	
Trough	The value shown is the minimum value found between the cursors measured relative to a baseline formed by joining the two points where the cursors cross the data. This is always less than or equal to 0.	
RMS error	This is the same as the SD except that the normalising factor is n , not $n-1$.	
Mean in X	This works for waveforms and result views. The value is the mean x value weighted by the data at each point. Although this would normally be used with positive data, it has some meaning with negative data values.	
SD in X	This is available for result views only. The result has a useful meaning only when all the data values in the time range are positive. The value is: $Sqrt(\text{Sum}(\text{data}[x] * (x-xm)^2) / \text{Sum}(\text{data}[x]))$ where the sum is over the x positions in the range and $\text{data}[x]$ is the data value at x and xm is the value returned by Mean in X.	
Mean Abs	This is the mean of the absolute values of the data values in the time range.	

Event channels After version 5.06, event channels drawn in mean and instantaneous frequency mode are treated the same as waveforms. For events drawn in other modes, all measurement types except those listed below produce a blank field in the window. There is an **Edit** menu Preferences option in the Compatibility tab to force the old behaviour, in which only the following measurement types produced a result:

Mean	The count of events or markers between the cursors divided by the time difference between the cursors.
Area	The total number of events or markers on the channel between the cursors.
Area (scaled)	The same as Area, but if a zero region is specified, the amount subtracted from the other regions is scaled by the relative width of the regions.
Sum	The same as Area.

Selecting, copying and printing data

Fields from this window can be copied to the clipboard or directly to the Log window. To do this, select the region to be copied and then click the right mouse in the window and use the **Copy** or **Log** command. To select an entire row or column, click the mouse in the titles at the top and left of the window. To extend a selection, hold down the Shift key. To select non-contiguous rows or columns hold down the control key and select the rows and columns as required. You can also print the entire window with **Ctrl+P** and set the font with **Ctrl+F**.

Horizontal cursors In many respects, horizontal cursors are similar to vertical cursors. However, they differ significantly in that each horizontal cursor is attached to a channel, or more accurately, to the y axis of a channel. If you change the drawing mode of an event channel with a horizontal cursor to a mode that has no y axis, the horizontal cursor will treat the available vertical space as if it runs from 0 at the bottom to 1 at the top.

If you drag channels with horizontal cursors on top of each other, only the horizontal cursor for the topmost channels (i.e. the channel with a valid y axis) is visible. In the special case of channels drawn with a locked y axis and a group offset of 0, all horizontal cursors for all channels are visible because the y axis is valid for all the channels.

New Horizontal



The command is available when a time or result window is the current window, and there are less than nine horizontal cursors already active. A new horizontal cursor is added to the top-most visible channel with a y axis. The cursor is given the lowest available horizontal cursor number and is labelled with the cursor label style for the window.

Delete Horizontal

The delete command activates a pop-up menu from which you can select a horizontal cursor to remove, or you can delete all the horizontal cursors. The cursors are listed with their number, channel number and position as an aid to identification. Deleting a cursor removes it from the window; other cursors are not affected.

Fetch Horizontal

This activates a drop down list from which you can select a horizontal cursor to place in the centre of the visible y axis range of the channel. This does not make a channel visible if the horizontal cursor is attached to an invisible channel.

Move To Horizontal

This command opens a pop-up menu from which you can select a horizontal cursor to move to. The cursors are listed with their number and position as an aid to identification. The window is scrolled vertically to display the nominated cursor in the centre of the y axis range.

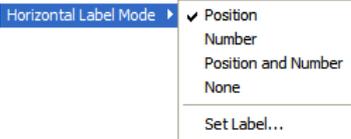
Position Horizontal

This command opens a pop-up menu in which you can choose a horizontal cursor and then opens a dialog in which you can set the position and channel for the cursor. You can also open the dialog by right clicking on a horizontal cursor and selecting **Set Position** from the cursor context menu.

Display all Horizontal

This has no effect if there are no horizontal cursors. If there is a single cursor, the command behaves as though you had used the Move To Horizontal command and selected it. When there are multiple cursors, the channel y axis range is adjusted such that the lowest cursor is at the bottom edge of the channel area and the highest is at the top edge.

Horizontal Label mode



Each cursor has an optional label used to identify it. You can drag the cursor labels to the left and right with the mouse to suit the data. There are five cursor label modes: **None**, **Position**, **Number**, **Position and Number**, and **User-defined**. You select the most appropriate for your application using the pop-up menu or by right clicking on a cursor and choosing to set the cursor label from the context menu. To avoid confusion between the cursor number and the position, the number is displayed in bold type when it appears alone and bracketed with the position.

Each cursor stores its own mode and label, and the view has a mode that is applied to new cursors. The first four items in the menu set the view mode and the mode of all cursors. You can also set a user-defined label for cursors (but not for the view) with the **Set Label** command.

Renumber Horizontal

When you create a horizontal cursor, they take the lowest available cursor number and you can also drag cursors over each other. In some circumstances you want the cursors numbered in order on the screen, for example, when taking the differences of cursor values. This command rennumbers the cursors, with cursor 1 at the bottom. Where cursors share the same channel, they are ordered so that the lower numbered cursor has the lower y axis value.

Cursor context commands

If you right-click on a horizontal or vertical cursor, a context menu opens. This holds commands that can be applied to items in the view that are close to the mouse position. In particular, it holds the item **Cursor n** for vertical cursor n or **HCursor n** for horizontal cursor n. Move the mouse pointer over one of these commands to open a new context menu.

This menu contains commands to set the cursor position, set the cursor label, copy the cursor value to the clipboard and delete the cursor. For vertical cursors, you can open the Active cursor mode dialog for that cursor. Horizontal cursors also have the option to copy the cursor position relative to any of the other horizontal cursors.

Sample menu

The **Sample** menu contains commands to configure the system for data sampling and to control the system during sampling.

Sampling configuration

This command opens the **Sampling Configuration** dialog that sets the sampling parameters used when you select the **File** menu **New** command. You can load and save the configuration from the **File** menu **Save Configuration** and **Load Configuration** commands. The dialog is described in detail in the *Sampling data* chapter.

Clear configuration

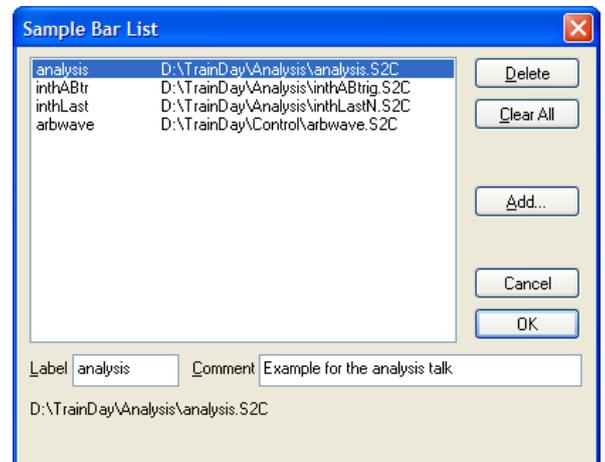
This command deletes all sampling window positions, associated result views for on-line processing, display trigger settings and WaveMark templates. It preserves the channel lists and all values visible in the **Sampling Configuration** dialog.

Sample Bar

You can show and hide the **Sample Bar** and manage the **Sample Bar** contents from the **Sample** menu. The **Sample Bar** is a dockable toolbar with up to 20 user-defined buttons. Each button is linked to a Spike2 configuration file. When you click a button, the associated configuration file is loaded and a new data file is opened, ready for sampling. You can also show and hide the **Sample Bar** by clicking the right mouse button on any Spike2 toolbar or on the Spike2 background.



The **Sample** menu **Sample Bar List...** command opens the **Sample List** dialog that controls the contents of the **Sample Bar**. Click **Add** to open a file dialog to select Spike2 configuration files (*.S2C) to add to the bar. If a file does not hold a label or comment, the first 8 characters of the file name form the label and the comment is blank.

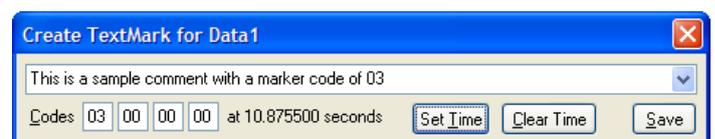


Select an item in the list to edit the label and/or comment for display; this does not change the configuration file. To re-order

Sample Bar buttons, drag them in the list. **Delete** removes the currently selected item. **Clear All** deletes all items. Spike2 saves the list in the registry. Each logon to Windows has its own registry settings, so each user has their own **Sample Bar** settings.

Create a TextMark (Ctrl+T)

This menu command opens the **Create TextMark** window. To use this command you must enable a



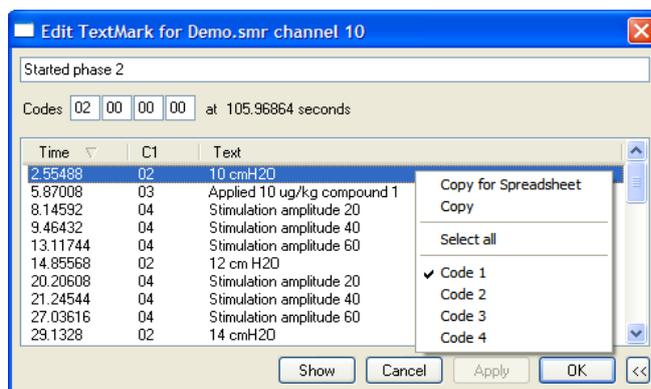
TextMark channel in the **Sampling Configuration**. You can always add markers by hand even if you enabled serial line input of **TextMark** data in the **Sampling Configuration**. The text of the last 10 markers is saved in the drop down list to save time when you have a repeating protocol.

You commit the comment to the file when you press the **Save** button. The **Set Time** button fixes the time associated with the comment, the **Clear Time** button clears the time. If you have not set a time, the time of the comment is the time when it was saved. If this channel is a trigger for Triggered sampling mode in the sampling configuration dialog, the trigger happens when you click the **Save** button and the **Save Time** and **Clear Time** buttons are replaced by the message **This channel is a trigger**.

- Comments are shown in the file as a small rectangle. The colour of the rectangle depends on the first marker code associated with the marker. If the code is 00, the rectangle is yellow. Otherwise, codes 01-08 use the same colours as WaveMark data, set by the **View** menu **Change Colours** command. The colours repeat every 8 codes.

You can see the text associated with a TextMark by moving the mouse pointer over the rectangle and waiting a moment. The text appears as a tool tip. Any movement of the mouse pointer hides the text. To edit the comment text and marker codes, double click the rectangle to open the **Edit TextMark** dialog. You can also use this dialog to navigate to any marker or range of markers in the file. The **Apply** button saves any changes without closing the dialog. **OK** saves changes and closes the dialog.

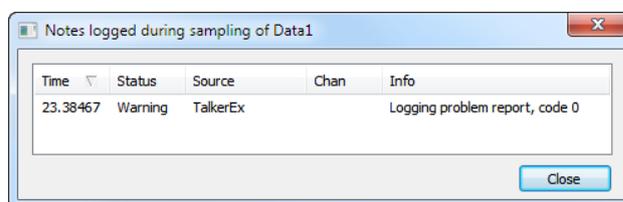
The list area of the dialog is hidden with the << button and restored with the >> button. Double-click an item in the list to edit it and save any changes to the item that was previously displayed. The **Show** button moves the time view display to the currently selected TextMark. If there are multiple selections, the first marker is displayed at the left edge of the screen and the last one at the right edge.



Right-click in the list to control how many marker codes to display and to copy selected data to the clipboard. If you choose **Copy for spreadsheet**, text strings and marker codes are enclosed in quotation marks so that the data will import easily into spreadsheet programs. You can change the sort order of the data in the list by clicking the column titles. Initially, the items are sorted based on the **Time** column.

Sampling Notes

Spike2 keeps the warning messages generated during sampling and this command displays them. The list of messages clears each time sampling starts. The count of notes is also displayed in the **Sampling Status** bar. You can right-click on the dialog to open a context menu from which you can copy the current selection to the clipboard.



Talkers

This menu command displays a list of Talkers that Spike2 knows about. You can select individual talkers to get Talker information, disable Talkers and to instruct Spike2 to forget about Talkers that are no longer required.

Offline waveform output

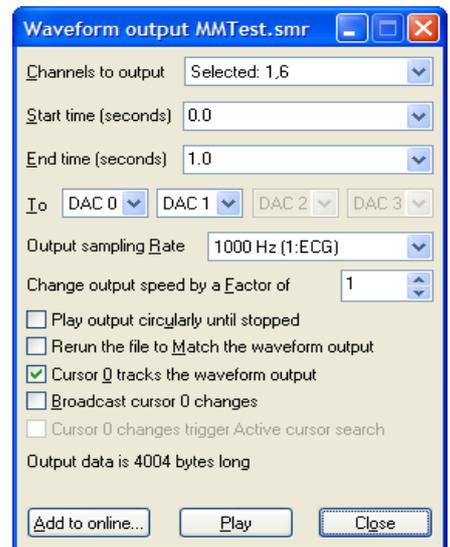
When a data file is active, the **Sample menu Waveform Output...** command opens the Waveform output dialog.

Channels to output

You can play up to 4 waveform or WaveMark channels through the 1401 DACs, or up to two channels through a sound card. To output multiple channels, select them in the time view, otherwise you can select single channels. You can change the channel selections after you open this dialog. Spike2 fills gaps in waveform data and gaps between WaveMark data items with zeros.

Time range

The **Start time** and **End time** fields set the time range to play. You can link these to cursors in the time view. The values used are those set when you click **Play**.

**To**

You can select either a 1401 DAC for each channel, or if you select **Sound** in the first drop down list, the first (and second) channel is played through the sound card (mono for one channel, stereo for two channels).

Output sampling Rate

You can choose the sampling rate of any of the channels for output, and the output will be played at a rate as close to this as the output hardware will allow. All channels are replayed at the same speed. If any channel has a different sample rate, Spike2 resamples the data to have the same rate as set by this field (using linear interpolation).

Change output speed by a Factor of

You can speed up or slow down output by up to 4 times. This changes the pitch of sound or the repeat rate of a repetitive signal. The actual rate depends on the output hardware.

Play output circularly

If you do not check this box, the waveform plays once only.

Rerun the file to Match the waveform output

If this box is checked when you click **Play**, the file will Rerun (simulate sampling) with the file time taken from the play position. This can be very useful when using Spike2 in teaching or for presentations, especially when using sound card output.

Cursor 0 tracks the waveform output

Check this box to have cursor 0 track the current playing position. Beware that cursor 0 has other uses in Spike2, for example in active cursors, and these may interfere with tracking the waveform replay position.

Broadcast cursor 0 changes

Check the *Cursor 0 tracks the waveform output* box, to enable this field. If you check it, each change made to cursor 0 for tracking purposes is broadcast to all linked dialogs and to any multimedia windows associated with the time view. This will update any open **Cursor Values** or **Cursor Regions** dialog, which may take appreciable time.

Cursor 0 changes trigger Active cursor searches

If the *Broadcast cursor 0 changes* box is enabled and checked, this field is enabled. Each change to cursor 0 will cause an active cursor search.

Play and Stop

This button starts output (and changes to **Stop**, which stops it). The output pays no attention to the scale and offset values associated with the channel. The 16-bit data values recorded for the file are played. If you have an 8-bit sound card and try to play a low level signal, the result may be disappointing (or just silence).

Add to online...

This option adds the selected waveforms to the on-line waveform output list. The data file must have been saved to disk. As long as the list of waveforms is not full and the current waveform is less than 32 MB in size you can **Add** the waveform to the list. You

can also **Replace** an existing waveform in the list. You can give each waveform a short name that labels the buttons for interactive replay of waveforms during data acquisition. There are more controls for the list of waveforms for replay in the Sampling Configuration dialog.

Waveforms can be saved as either a data file, time range and a list of channels or by converting the current channel list into the memory image that would be played through the output DACs on line. If a waveform has been converted the Source column in the list starts with “Memory”. The advantages of converting are:

- The output is not affected if the original file is moved or deleted or changed
- The (usually small) time for the conversion is saved each time you sample data

However, there are also disadvantages to converting:

- It can take a lot of memory to hold the converted data which can slow Spike2 down and may cause your system to run out of memory
- If the sampling configuration is saved, the converted data is also saved, which makes the configuration files much larger and can slow down system operation.

To prevent the system running out of memory, there is a limit on the amount of memory (32 MB) that any one converted wave may use. If any channel in the play list is a WaveMark or a memory channel or a duplicate channel, Spike2 will always attempt to convert the data, as there is no guarantee that these channels will exist (or have the same marker filter) when sampling is requested.

If a channel has a channel process applied to it, you must either leave the data file open with the channel process applied or you must convert the wave. This is because a channel process only has effect while the data file is open in Spike2.

Sampling controls

This command hides and shows the Sample control toolbar. The controls within this toolbar are duplicated in this menu as the **Start sampling**, **Write to disk**, **Abort sampling** and **Reset sampling** commands (see the *Sampling data* chapter for details of the floating window and menu commands).

Sequencer controls

This command hides and shows the sequencer control panel that is available during sampling (see the *Data output during sampling* chapter for details).

Change Output Sequence

This command is used during data sampling to replaces the current output sequence. It displays a File Open dialog in which you can select an output sequence file. See the *Data output during sampling* chapter for more information.

Conditioner Settings

Spike2 supports serial line controlled programmable signal conditioners. These devices amplify and filter waveform signals, and can provide other specialist functions. If a suitable conditioner is installed in your system, this menu command opens the conditioner dialog (see the *Programmable signal conditioners* chapter for details).

Discriminator Configuration

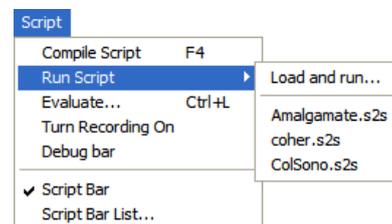
If your 1401*plus* has a CED 1401-18 event discriminator fitted, this command opens a configuration dialog (see the *1401-18 Programmable Discriminator* chapter for details).

Script menu

The **Script** menu gives you access to the scripting system. From it you can compile a script, run a loaded script, evaluate a script command for immediate execution, record your actions as a script and convert a script from the format used in the DOS version of Spike2 to this version. You can find details of the script language and a description of the script window in the separate manual *The Spike2 script language* and in the on-line help. The menu commands are:

Compile Script This option is enabled when the current window holds a script. It is equivalent to the compile button in the script window. Spike2 checks the syntax of the script, and if it is correct, it generates a compiled version of the script, ready to run.

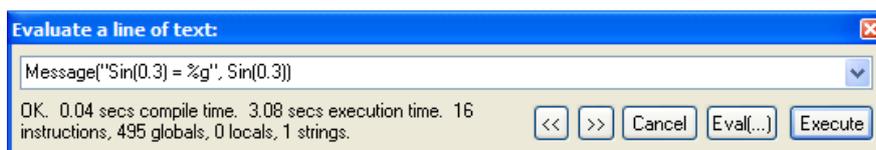
Run Script This option pops-up a list of all the scripts that have been loaded so that you can select a script to run. Spike2 compiles the selected script and if there are no errors, the script runs. If you run a script twice in succession, Spike2 compiles it for the first run, and uses the compiled result for the second run, saving the compilation time. If a script stops with a run time error, the script window is brought to the front and the offending line is highlighted.



You can also select the **Load and Run...** option from which you can select a script to run. The script is hidden and run immediately (unless a syntax error is found in it).

Evaluate

This command opens the evaluate window where you can type in and run one line scripts. You can cycle round the last 10 script lines with the << and >> buttons. The **Execute** button compiles and runs your script, **Eval()** does the same, and also displays the value of the last expression on the line.



Turn Recording On/Off You can record your actions into the equivalent script. Use this command to turn recording on and off. When you turn recording off, a new script window opens that holds the commands that are equivalent to your actions. When recording is enabled **REC** appears in the status bar.

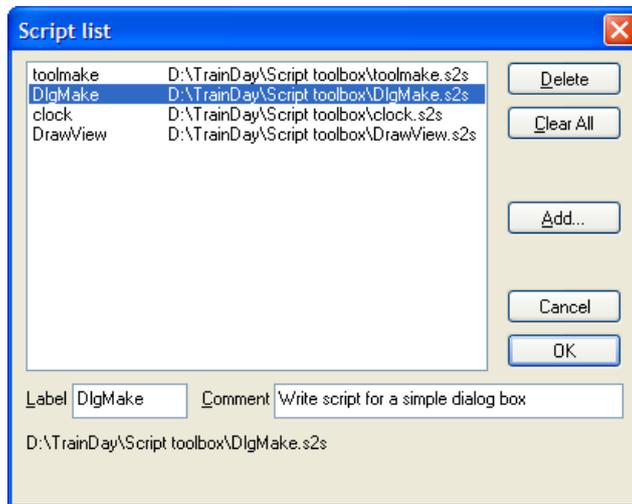
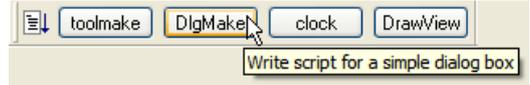
Debug bar You can show and hide the debug bar from this menu when the current view is a script. You can also show and hide the debug bar by right clicking on the title bar of the Spike2 window, or on the application window.

Convert DOS Script

This option is only available when a text window is open that is not empty. The command will attempt to convert a Spike2 for DOS script into the equivalent commands for this version of Spike2. This option is described in detail in the script manual.

Script Bar

You can show and hide the Script Bar and manage the Script Bar contents from the Script menu. You can also show and hide the Script Bar by clicking the right mouse button on any Spike2 toolbar or on the Spike2 background. The Script Bar is a dockable toolbar with up to 20 user-defined buttons. Each button is linked to a Spike2 script file. When you click a button, the associated script file is loaded and run. There is also a user-defined comment associated with each button that appears as a tool-tip when the mouse pointer lingers over a button.



The Script menu Script Bar List... command opens the Script List dialog from where you can control the contents of the Script Bar.

The Add buttons opens a file dialog in which you can choose one or more Spike2 script files (*.S2S) to add to the bar. If the first line of a script starts with a single quote followed by a dollar sign, the rest of the line is interpreted as a label and a comment, otherwise the first 8 characters of the file

name form the label and the comment is blank. The label is separated from the comment by a vertical bar. The label can be up to 8 characters long and the comment up to 80 characters. A typical first line might be:

```
'$DlgMake|Write script for a simple dialog box
```

You can select an item in the list and edit the label and comment. This does not change the contents of the script file. You can re-order buttons in the bar by dragging items in the list. The Delete button removes the selected item. Clear All removes all items from the list.

The list of files in the Script Bar is saved in the registry when Spike2 closes and is loaded when Spike2 opens. Each different logon to Windows has a different configuration in the registry, so if your system has three different users each has their own Script Bar settings. Alternatively, you can have different experimental configurations by logging on as a different user name.

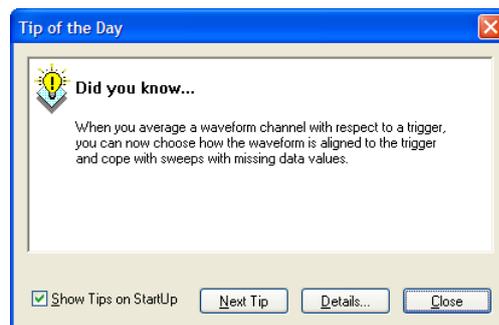
Help menu

Using help Spike2 supports context sensitive help that duplicates the contents of this manual and the script language manual. You can activate context sensitive help with the F1 key from most dialogs to get a description of each field in the dialog. You can use the help Search dialog to lookup topics that are not covered by the index.

In the script window you can obtain help by placing the cursor on any keyword in the script and pressing F1. To get help on a script function, type the function name followed by a left hand bracket, for example `MemChan (`, then make sure that the cursor lies to the left of the bracket and in the function name and press F1.

The help system has indexes, hypertext links, keyword searches, help history, bookmarks and annotations. If you are unsure of using help, once you have opened the help window, use the Help menu Using help command for detailed instructions.

Tip of the Day You can show the Tip of the Day window when Spike2 starts by checking the box. Of course, if you hate this sort of thing, clear the check box and you never need see it again! The Details... button opens the on-line Help at a page that holds more information about the topic in the window. The tips are held in the hidden file `sonview.tip` in the same folder as Spike2. If you have a useful tip that you would like to share with others, send it to us and we will very likely include it in the next release.



View Web site If you have an Internet browser installed in your system, this command will launch it and attempt to connect to the CED web site (www.ced.co.uk). The site contains downloadable scripts, updates to Spike2 and information about CED products.

Getting started This command runs a demonstration script that gives a quick tour of Spike2 basics. The option is disabled if the script is not present or if you are sampling data.

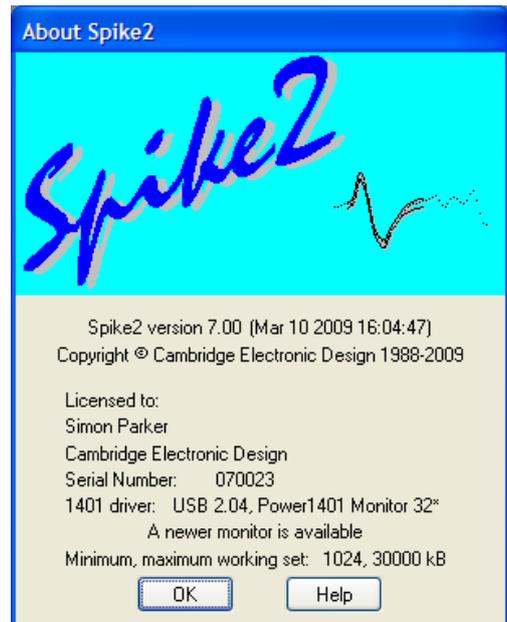
Other sources of help If you are having trouble using Spike2, please do the following before contacting the CED Software Help Desk:

1. Read about the topic in the manual. Use the Index to search for related topics.
2. Try the help system for more information. Use Search to find related topics.
3. If the problem is in sampling data, run the Try1401 program provided as part of Spike2. This checks out the 1401, device driver and interface card.
4. If none of the above helps, FAX, email (softhelp@ced.co.uk) or call the CED Software Help Desk (see the front of this manual for details). Please include a problem description, Spike2 serial number and program version and a description of the circumstances leading to the problem. It would also help us to know the type of computer you use, how much memory it has and which version of your operating system you are running.

About Spike2 This command is found in the Help menu. It opens an information dialog that contains the serial number of your licensed copy of Spike2, plus your name and organisation. Please quote the serial number if you call us for software assistance.

1401 device driver If there is a 1401 device driver installed, the driver revision is displayed. If the driver is older than Spike2 expects, you will be warned. Spike2 displays the type of 1401 and the monitor version if a 1401 is connected and powered up.

1401 Monitor revision If the monitor is not the most recent at the time this version of Spike2 was released, an asterisk follows the version. If it is so old that it compromises data sampling, two asterisks follow the version.



The Power1401 and Micro1401 have firmware in flash memory. Flash updates and instructions for applying them are available as downloads from the CED web site; you can update the flash firmware without opening the 1401 case. These updates contain bug fixes, code improvements and sometimes new hardware features.

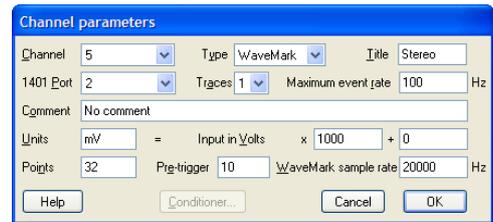
New monitor ROMs are available from CED for the 1401*plus* and the micro1401. You must open the 1401 case to replace them; we ship detailed instructions with the ROM.

WARNING: 1401 monitor too old or interface problem This message appears if we detect that the 1401 monitor ROM is too old to run Spike2 safely or if there is a problem with 1401 communications such that the 1401 is detected but communications are garbled. If the monitor ROM is OK, this can be caused by a 1401 interface card having an address or interrupt vector problem. Contact CED for advice.

Working set size If you are running Windows XP or a later operating system, there is information about the Working Set Size at the bottom of the About box. The two numbers describe the minimum and maximum physical memory that the operating system allows Spike2 to use. If you suffer from error -544 when you sample data, these numbers are important. You will find more detailed information in the on-line help (look up “Working Set” in the on-line help index).

Sampling parameters

To sample spike shapes you must create a WaveMark channel in the Sampling configuration dialog. The Title, Port, Channel comment, Units and Scaling fields are the same as for Waveform data. The Maximum event rate field is the total number of spikes (of all classes) that you expect per second on this channel.



The Traces field sets the number of waveforms that are combined in this channel. This should be set to 1 for a single electrode, 2 for a stereotrode and 4 for a tetrode. The 1401 Port sets the input channel number of the first trace. The remaining traces are taken from sequential channel numbers. For example, if the 1401 Port field was set to 7 and Traces was 4, the 1401 input channel numbers 7, 8, 9 and 10 would be used for the traces.

The bottom line of the dialog box sets the number of points to save per trace per spike in the range 10 to 126 (Points) and the number of points to display before the peak (Pre-trigger). You can alter the number of points before the peak and the total number of points in the template setup section, so it is not essential to get this exactly right.

The WaveMark sample rate field sets the ideal sample rate for all WaveMark channels. A sampling rate of 20 kHz per channel is often used so that spikes (of typical duration 1 to 2 milliseconds) can be usefully discriminated. The actual sampling rate depends on the number of waveform and WaveMark channels. See the Sampling configuration Resolution tab for more information on setting waveform and WaveMark sample rates.

Template matched signal

You can use digital output bits 0-7 (on the rear panel of Micro1401 and Power 1401) to signal that a spike has matched a template. There are two output formats (set in the on-line Template Setup). In One bit mode, a single bit pulses from low to high, then back to low for each of up to 8 templates. Bit 0 is for the first template, bit 1 for the second, and so on up to bit 7 for the eighth template. One bit mode is not useful when you have more than one channel of spikes as you cannot tell which channel produced the output.

In Coded mode digital output port pin 23 pulses for each matched template, and digital output bits 3-7 hold the bottom 5 bits of the ADC port number and bits 0-2 hold the template (0-7 for the 8 templates). The data is valid on both edges of the pulse. The pulse is short, less than 1 μ s for the Power1401, more than 1 μ s for other 1401s.

Digital output connections	Data bit	7	6	5	4	3	2	1	0	Gnd	Strobe
	Output pin	5	18	6	19	7	20	8	21	13	23

1401*plus* digital bits 0-7 are used for both input and output. They are normally inputs. The state of bits used as outputs read back in the sequencer DIGIN and WAIT instructions. Coded mode uses all 8 bits, One bit mode uses one bit per template. The Micro1401 and Power1401 have independent input and output bits 0-7, however you should use the output sequencer DIGLOW command with caution as it controls the same output bits.

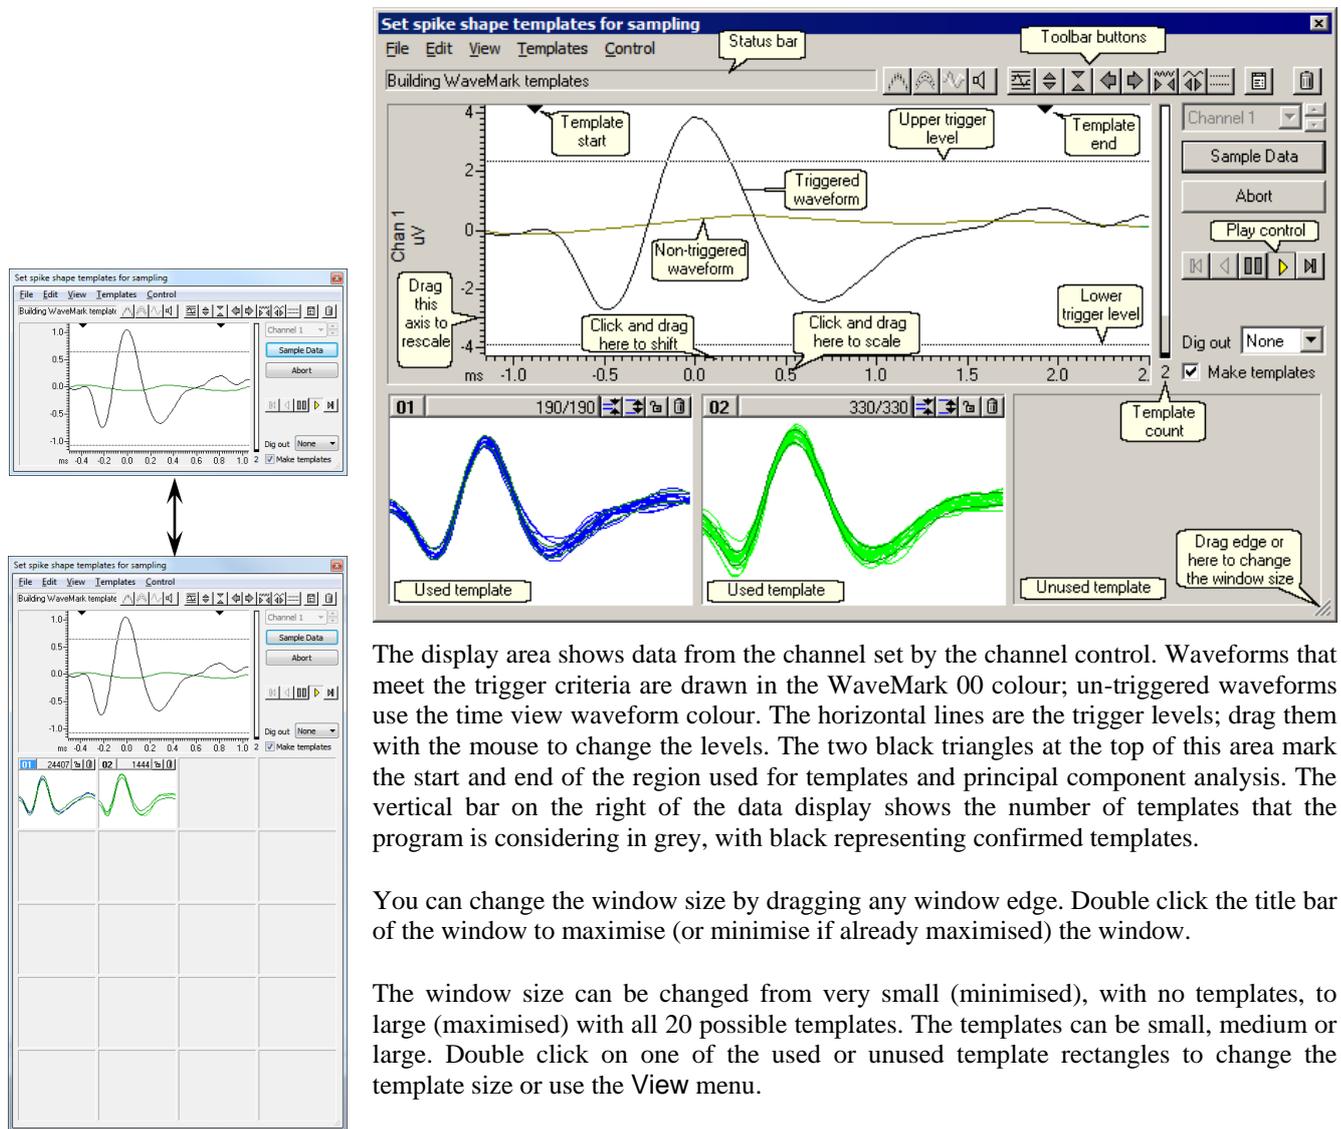
Timing of the digital output

Measurements with templates of 20 points show that the vast majority of spikes are matched within 2 milliseconds of the end of the spike (much less than this with a Power1401). However you cannot rely on this. Even with a PCI interface card, and a fast host, delays of 30 milliseconds or more may occur occasionally, especially when data is written to disk at high sampling rates. Longer delays will occur with slower hardware.

You can test the delays by feeding the output back into an event channel and correlating the Spike channel with the event channel. If you intend to use these outputs for a timing critical use it is important that you measure the typical time delays for your configuration.

On-line template setup

When you create a New file with WaveMark channels, Spike2 opens the template setup window. The window has its own menu and contains four main regions: a status and toolbar at the top, a data display area with the trigger levels, a control area on the right and a region with up to 20 templates at the bottom. The first 8 templates are used on-line.



Channel control

The display area shows data from the channel set by the channel control. Waveforms that meet the trigger criteria are drawn in the WaveMark 00 colour; un-triggered waveforms use the time view waveform colour. The horizontal lines are the trigger levels; drag them with the mouse to change the levels. The two black triangles at the top of this area mark the start and end of the region used for templates and principal component analysis. The vertical bar on the right of the data display shows the number of templates that the program is considering in grey, with black representing confirmed templates.

You can change the window size by dragging any window edge. Double click the title bar of the window to maximise (or minimise if already maximised) the window.

The window size can be changed from very small (minimised), with no templates, to large (maximised) with all 20 possible templates. The templates can be small, medium or large. Double click on one of the used or unused template rectangles to change the template size or use the View menu.

When there is more than one channel holding WaveMark data, you can change the channel either by dropping down the list of channels and selecting one, or by using the small spin box on the right of the channel control to move to the next channel in the list. Each channel has its own set of templates and template parameters.

Sample data Once you have adjusted the trigger levels and set the templates for each channel, click this button to close the window and sample data.

Abort This button closes the window and does not sample data.

Dig out During sampling, the 1401 can flag template matches using bits 0-7 of the digital port. You can select between **None** (no output), **One bit** (single pulsed digital bit per template) or **Coded** (data channel and template coded in the 8 data bits).

Make templates With this box checked, detected spikes are offered to the template matching system and will create and modify templates. With this box unchecked, spikes are compared against

templates but the templates do not change. When you open this window, this box is checked if there are no existing templates, otherwise it is clear. This command is also available in the **Templates** menu with a keyboard short-cut of `Ctrl+M`.



These five buttons control data replay. The centre button pauses replay and the buttons either side of it play the data forwards and backwards. The outer buttons step one spike forwards and backwards. You can only play or step backwards in the **Edit WaveMark** command. There are keyboard shortcuts for these controls. `B` steps backwards one spike, `M` steps forwards one spike, `P` runs forwards, `Q` runs backwards and `V` pauses output. You can also access the **Play** commands from the **Control** menu.

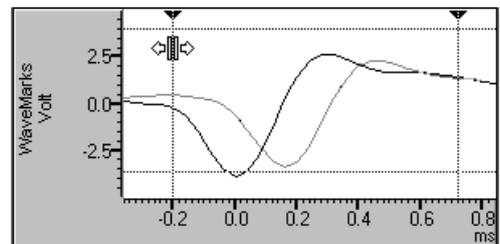
Renumber... This command is available from the **Templates** menu with a keyboard short-cut of `Ctrl+R`. It prompts you for a starting marker code then renumbers the templates with consecutive codes, starting with the code you set. The templates remain in the original positions and code order. If two or more templates share a marker code they will still share a code after renumbering. For example, templates coded 01, 04, 04, 07 and 03 renumbered to start at 01 would end up with codes 01, 03, 03, 04 and 02.

Order by code This command is available from the **Templates** menu with a keyboard short-cut of `Ctrl+B`. It rearranges the templates so that they are in ascending marker code order. For example, templates with marker codes 01, 04, 04, 07 and 03 would be sorted into the order 01, 03, 04, 04 and 07. The sorting order is left to right and top to bottom.

Load and Save... This command is available from the **File** menu with a keyboard short-cut of `Ctrl+S`. It opens a new dialog in which you can load templates generated in previous spike sorting sessions and save your templates to resource files.

Selecting the area for a template

The two black triangles at the top of the display area set the waveform region used for template formation and principal component analysis. To adjust the region, click the triangles and drag. Vertical cursors appear as guides for the new template position and the cursor positions and template size and start position relative to the trigger point appear in the status bar. Old templates are deleted if you change the number of points in the template.



In most cases the best information about the spike class is found around the peaks; the start and end of the spike record holds mainly baseline noise. If you use the entire spike shape for template matching, you will be comparing baseline noise for some of the record, which wastes time and reduces the sensitivity of the template matching.

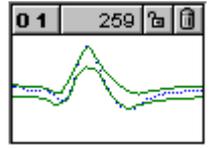
If your spike only fills a small proportion of the spike trace you should reduce the number of points for the **WaveMark** channel with the  button. This has several advantages: reduced disk space for storage, faster processing, quicker drawing of data and fewer records with a second spike overlapping the end.

If you have too few points around the interesting region of your spike you will need to increase the sampling rate. The template routines interpolate between data points, so you only require enough points to define the spike shape.

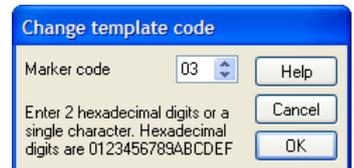
You will get the best spike discrimination when the area selected shows the greatest variation between different classes of spikes. For example, if all your spike classes look very similar at the end, you will improve discrimination by excluding the end of the data.

The template area

Templates appear at the bottom of the window. If this area is not visible, drag the bottom edge down. Templates can be small, large or huge; double click a template to change the size. The template code is at the top left of the box; when a template first appears, it has the lowest unused code. The count of spikes that matched the template is in the centre. In large and huge mode, the display also shows the number of spikes added to the template as: matched/added. You can set the count of matched events to 0 with the Template menu Initialise Counters command; the short-cut key is `Ctrl+I`.



Double click the template code to open a dialog where you can change the code. Double-digit codes are read as hexadecimal marker codes, single characters are read as printing characters. If you have a spike that changes shape too much for one template to represent it, let it generate several templates and then edit the codes to be the same. You can use the code 00 to mark spikes or artefacts you are not interested in.



The display area holds the last spike that matched the template, the template itself and optionally all non-matching spikes. You control the drawing style of the matching and non-matching spikes and the template with buttons at the left-hand end of the toolbar. These buttons can appear at the top of each template area:

Change width 

These buttons are hidden in small template windows. They decrease and increase the width of the template.

Lock 

This button locks or unlocks a template. A locked template does not change when new spikes match it. An unlocked template adapts as each new spike is added. Templates can lock automatically if Auto Fix mode is set in the template settings dialog. When building templates, Spike2 cannot merge a provisional template with a locked template. Instead, Spike2 creates a new template with the same code as the one it would have merged with.

Clear 

This button erases the template. To erase all templates, click the bin icon in the toolbar.

Merging templates

You may decide that two templates are so similar that you would like to merge them into one. To do this, drag one of the templates over the one that you wish to merge it with and release the mouse button. The dragged template vanishes and the template over which you released it changes to take account of the new data.



The mouse pointer is an open hand when it is over a spike or template that you can drag. To drag, click and hold the left mouse button; the pointer changes to a closed hand. When you drag over a drop zone, the closed hand has a plus sign on the back. You can compare templates by dragging, but not releasing the mouse. If you decide not to merge, make sure the mouse pointer is not the closed hand with the plus sign when you release the mouse button.

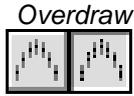
Manual template creation

When paused, you can use drag and drop to create new templates and to classify spikes. Click the mouse in the data display area and drag the current spike to an existing template to set the marker code, or to an unused template to create a new template.

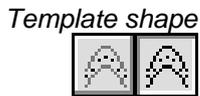
You can also use the keyboard to do the same thing: The keys 1 to 9 will match the current spike to the first template with the code 01 to 09, or if there is no template with a matching code, a new template will be created with that code. The 0 key can be used when editing spikes to give the current spike the code 00.

Toolbar controls

The toolbar buttons control the template forming process. The leftmost group are toggles; click them to set, click again to clear. The remaining buttons act immediately. These commands are also available in the dialog menus and most have short-cut keys:



Normally, the templates display the last spike added. You can choose to show all spikes added to templates by overdrawing. This omits the step of erasing the previous spike before drawing a new one. The short-cut key is `Ctrl+O`.



You can choose to display the mean template, or the upper and lower limits of the template. The template limits are useful when you display the last spike and need to see how well it fits within the template. The mean template can be useful when you are comparing template shapes. The short-cut key is `Ctrl+L`.



In some circumstances you need to see every spike displayed over every template. With this button down, all non-matching spikes are drawn on top of the template in the "Not saving to disk" colour. The short-cut key is `Ctrl+N`.

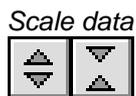


If you have a sound card, turn this on to hear a tone for each spike. A different tone is played for each template code; the tone is not related to the spike shape. A low-pitched tone is played for a spike that matches no template, and a high-pitched tone for spikes that match templates numbered higher than 20. The short-cut key is `Ctrl+U`.

There are limits on the maximum spike rate imposed by the operating system and the fact that sounds have to last much longer than the spikes for you to hear the tones!



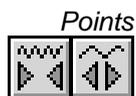
This uses information collected from recent spikes to set the trigger and limit levels and to scale the data. This gives a good starting point for manual adjustment. The display scaling applies to both the data window and to the templates. The short-cut key is `Home`.



These two buttons increase and decrease the display size of the waveforms and the templates. If you click on a button, the scale changes once. If you hold one down, the scale changes repeatedly until you release the button. You can also scale the data by clicking and dragging the y axis. The short-cut keys are `Up` arrow and `Down` arrow.



These buttons scroll the data in the large window sideways. This changes the pre-trigger time, and is equivalent to changing the `Pre-Trigger` field of the `Channel Parameters` dialog for `WaveMark` data. If possible, the two triangles marking the start and end of the template region also scroll to preserve the template region. You can also click and drag the data x axis ticks to scroll the window or use the `Left` and `Right` arrow keys.



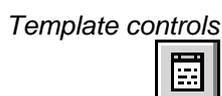
These buttons increase and decrease the number of points saved for each `WaveMark`. This is the same as the `Channel Parameters` dialog `Points` field. You can also use the `PgUp` and `PgDn` keys or click and drag the x axis numbers to scale the view around the 0 ms point. You should save as few points as possible to minimise the file size.



Depress this button to display and use the limit cursors. Spikes with peak amplitudes above the limits at the trigger point are ignored. The short-cut key is `Ctrl+H`. Online, limit cursors are available with a `Power1401` or a `Micro1401-3`. Offline, they are available when creating `WaveMark` data from a waveform.

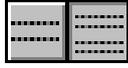


This button clears all templates (and any tentative templates not yet confirmed). Think of throwing the templates in the bin. The short-cut key is `Del`. You can clear individual templates using the clear button on the top bar of each template.



This button opens a dialog that controls template generation. To use the dialog you need to understand how templates are constructed and generated, (see page 15-8, *Template formation* for more information). The short-cut key is `Ctrl+Enter`.

Horizontal cursors



All Spike shape dialogs (except off-line Edit WaveMark) have a choice of 2 (setting upper and lower trigger levels) or 4 horizontal cursors (setting amplitude limits and trigger levels) controlled by a toolbar button. You can link the cursors to horizontal cursors (1=lower trigger, 2=upper trigger, 3=lower limit, 4=upper limit) in the associated time view source channel; when linked, dragging and releasing a cursor updates the linked cursor. Right-click a cursor or in the cursor area to open a context menu:

Set cursor position...



This opens a dialog in which you can enter a new position. **Reset** restores the original position. **Move away** closes the dialog and moves the cursor to a position where it will have no effect on spike capture. The drop down list selects horizontal cursor positions in any associated time view(s). **OK** is enabled if the new position is legal and different from the original.

Copy position

Select this option to copy the current cursor position and units to the clipboard as text.

Move away

Move the cursor to a position where it has no effect on spike triggering. If you have enabled the amplitude limit cursors and you use the this command on a trigger level, both the trigger and associated amplitude cursor will move away.

Fetch cursors

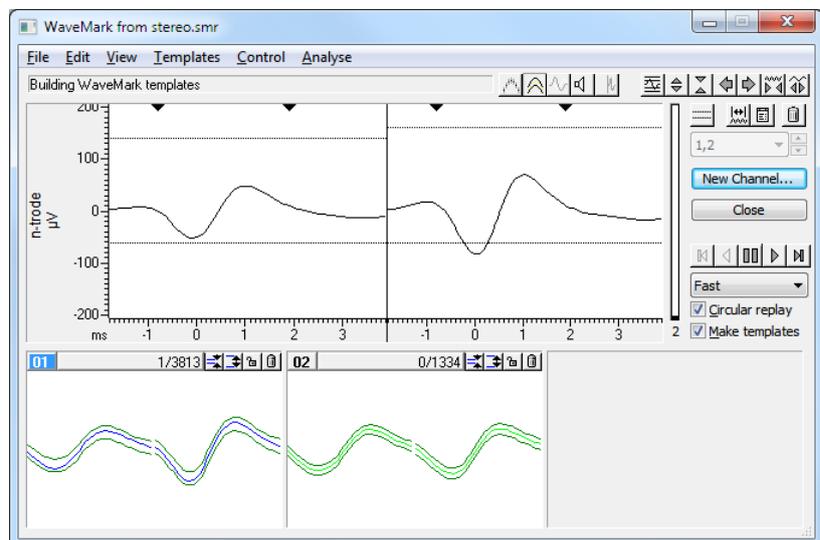
This appears if you right-click in the display area and a horizontal cursor is beyond the displayed range. The command moves the cursor back into a visible range. If there are multiple traces, only cursors in the trace on which you click are restored. If the cursors are linked to the time view, the fetched cursors are also linked.

Link Time view cursors

This links and unlinks the visible horizontal cursors in the dialog to horizontal cursors 1 to 4 in the time view channel that sources the data. This option is disabled if the Cursor menu is disabled during script operation. See the online Help for more.

Multiple traces

All the spike shape dialogs can handle WaveMark data with 1, 2 or 4 traces (single electrode, stereotrode and tetraode data). The documentation concentrates on the single trace case; this section explains the differences when you have multiple traces.



The data display shows each trace in its own rectangular area with independent horizontal cursors (when used). The x axis and the triangles that mark the start and end of the template region are duplicated for each trace and are slaved together; if you drag the x axis or the template markers in one trace, all traces change. The templates are also divided up horizontally, however you can choose to overdraw them (perhaps to compare latencies) with the View menu **Overlay templates** command.

Spike sorting and template matching work in much the same way for multiple traces as for a single trace. The template parameters apply to each trace and all traces must meet the matching criteria for a spike to match a template. Because of this, you may need to relax some of the matching criteria compared to spike sorting with a single trace.

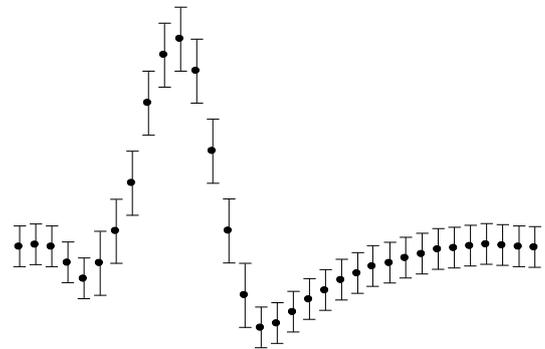
Forming n-trode data from waveform channels

You can create WaveMark data from 2 or 4 existing waveform channels with the same sampling rate. Select the channels in the time view by clicking the channel numbers, then use the Analysis menu **New n-trode** command or right click a selected channel and select **New Stereotrode** or **New Tetrode**. This will open the New WaveMark dialog with the selected channels set as the data source. The channel selector in the dialog displays the channels that correspond to the traces in the order that you selected them.

Template formation

A template is stored as a series of data points. Each point has a width and a minimum width associated with it. The width represents the expected error in a signal that matches the template at that point. Spikes match a template if more than a certain percentage of the points in a spike fall within the template. It is easy to calculate appropriate widths from a population of spikes of the same class. The problem comes when you create a new template from a single spike. What width do you give each point?

We set the initial template width to a percentage of the maximum positive or negative amplitude of the spike. The percentage to use depends on the data, 35% is a reasonable first guess. This value also sets the minimum allowed spike width. Unless a minimum width is set, the template width tends to reduce as we only add matching spikes. We also limit the sum of the differences between a spike and the mean template.



The template width is twice the mean distance between the template and the spikes that created it, but is not allowed to become less than the minimum. If the variation in spike shape around the mean were normally distributed, you would expect 80% of the data in a matching spike to lie within the template. The template building algorithm is:

- 1 The first spike forms the first template with the template width estimated. This is a *provisional* template. Provisional templates are not displayed.
- 2 Each new spike is compared against each template in turn to see if sufficient points fall within the template. If there are *confirmed* templates, these are considered first, and if a match is found, the provisional templates are not considered.
- 3 If a spike could only belong to one template, it is added to it. Otherwise, the spike is added to the template that is the minimum distance from the spike. Distance is defined as the sum of the differences between the spike and the template.
- 4 If a spike belongs to no templates, a new provisional template is created.
- 5 Each time a spike is added to a template, all provisional templates have their decay count reduced by 1. If any decay count reaches 0, that template has 1 spike removed from its spike count and the decay count is reset. If the number of spikes in a provisional template reaches 0, the provisional template is deleted.

Adding a spike to a template changes the template shape and width. Once a pre-set number of spikes have been added to a template, it is checked against existing confirmed templates to prevent generation of two templates for the same data. If no match is found, the template is promoted to confirmed. If a match is found, the provisional template is

merged with the confirmed one unless the confirmed template is locked, in which case a new confirmed template is created *with the same code as the one it matched*. Each time a spike is added to a provisional template, the provisional template *decay count* is reset.

Step 2 is complex. To compare a spike against a template we shift it up to 2 sample points in either direction to find the best alignment point. Shifts of a fraction of a point are done by interpolation. If amplitude scaling is enabled, the spike is scaled to make the area under the spike equal to the area under the template (limited by the maximum scaling allowed). Amplitude scaling is only allowed for confirmed templates.

Template settings dialog

The dialog controls the formation of new spike templates and how spike match templates. The buttons at the bottom of the dialog are:

Copy to All

Use this button when you have multiple channels of WaveMark data and you wish to apply the current template parameters to all of the channels. The action is the same as **Apply**, but it sets the value for all channels, not just for the current channel.

Apply and OK

These do the same thing, but **OK** closes the dialog and **Apply** leaves it open. The settings you edit in the dialog make no difference to the template formation process unless you use one of the **Copy to All**, **Apply** or **OK** buttons. You cannot use these buttons if any field in the dialog holds invalid data.

Cancel

This button closes the dialog without applying any changes. Any changes made by previous use of the **Apply** or **Copy to All** buttons are preserved.

Each channel has an independent set of parameters. If you use the **Channel** control in the parent spike shape dialog to change data channel, the template parameters will also change to match the values for the new channel and any changes made will be lost unless you have used the **Apply** or **Copy to All** buttons to save the values.

The dialog has 4 areas:

New Template

This controls the creation of new templates.

Number of spikes for a new template

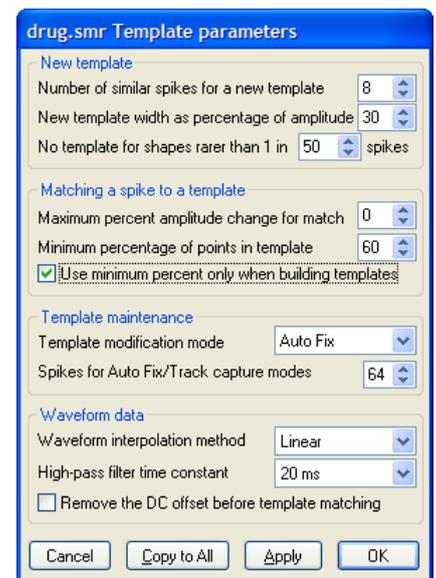
This sets the number of spikes at which a provisional template is promoted to a real one. We find that 8 is a reasonable starting value.

New template width as a percentage of amplitude

This is the percentage of the spike amplitude to give the initial (and minimum) template width. As spikes are added to the template, the width changes to represent the variation in amplitude of the spikes in the template. The maximum template width is set to 4 times the minimum width. A value of 30% is a reasonable starting value.

No template for shapes rarer than 1 in n spikes

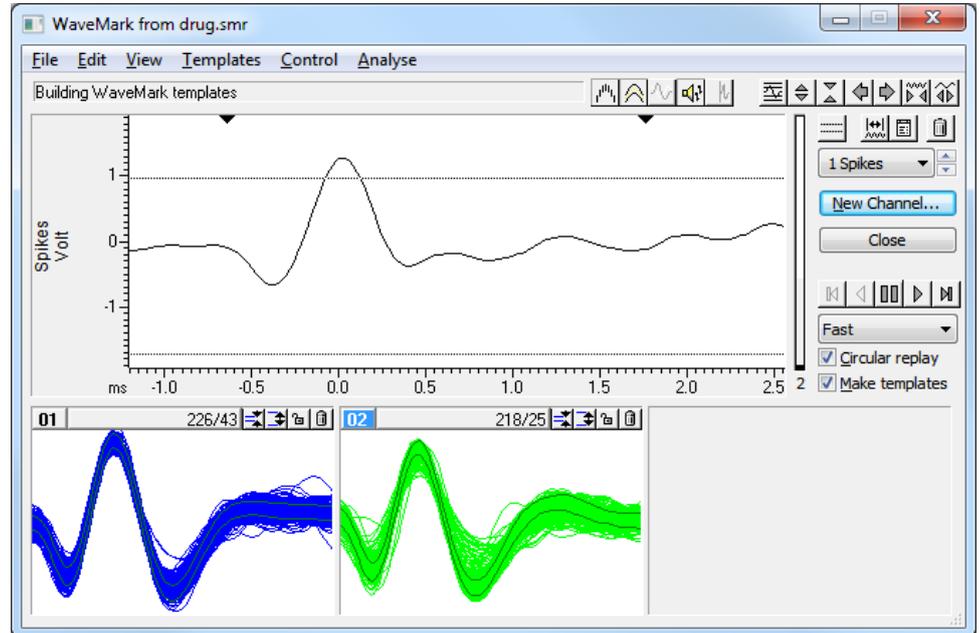
If this is n, this is roughly the same as saying that you are not interested in spike classes which occur less often than once every n total spikes. If you want to keep all spikes as potential templates you should set this to a large number. We find that 50 is a reasonable starting value.



Matching a spike to a template	The items in this group are used when comparing a new spike with the existing templates to determine if the new spike is the same or should start a new provisional shape. As well as the conditions mentioned below there is always a limit to the error between a spike and a template (unless amplitude scaling is enabled) to prevent totally ridiculous results and to avoid wasting time interpolating spike shapes for data that could never fit a template.						
<i>Maximum percent amplitude change for match</i>	Spike2 will scale spikes up or down by up to this percentage to make the area under the spike the same as the area under each target template. This is useful if you have spikes that maintain shape but change amplitude. Avoid using this on-line as it slows down the template matching process. Set this to 0% unless you need it. The maximum is 100%, which allows a spike to be doubled or halved in amplitude. The value you set depends on the amplitude variability of your spikes; 25% is a reasonable starting value.						
<i>Minimum percentage of points in template</i>	The percentage of a spike that must lie within a template for a match. If more than one template passes this test, spikes are matched to the template with the smallest error between the mean template and the (scaled) spike. 60% is a reasonable starting value. If you enable amplitude scaling you may want to increase this value to 70% or more.						
<i>Use minimum percentage only when building templates</i>	The template width is most useful in the setup phase when we are creating templates. Once templates are established it can sometimes be better to match to the template with the smallest error and ignore the width. If you check this field, this sets the percentage of points that must lie in the template to 0% unless you are building templates.						
Template maintenance	This group of fields determines how the shape of a template changes as more spikes are added to it. On-line, the template shape is fixed in all modes apart from Track. From our experience, Add All and Auto Fix are the most useful modes.						
<i>Template modification mode</i>	<table><tr><td>Add All</td><td>All spikes that fit the template are added and modify the template. The effect of each spike becomes smaller as the number of matched spikes gets larger.</td></tr><tr><td>Auto Fix</td><td>The template is fixed once a set number of spikes have been added. If you have several similar spikes, using Auto Fix after a fairly small number of spikes can stop a template gradually changing shape and becoming the same as another template.</td></tr><tr><td>Track</td><td>The template shape tracks the spikes. The contribution of each spike to the template decays as more spikes are added. This is only really useful for slow changes in spike shape and always brings with it the danger than all your shapes will merge together. It also slows down on-line spike classification.</td></tr></table>	Add All	All spikes that fit the template are added and modify the template. The effect of each spike becomes smaller as the number of matched spikes gets larger.	Auto Fix	The template is fixed once a set number of spikes have been added. If you have several similar spikes, using Auto Fix after a fairly small number of spikes can stop a template gradually changing shape and becoming the same as another template.	Track	The template shape tracks the spikes. The contribution of each spike to the template decays as more spikes are added. This is only really useful for slow changes in spike shape and always brings with it the danger than all your shapes will merge together. It also slows down on-line spike classification.
Add All	All spikes that fit the template are added and modify the template. The effect of each spike becomes smaller as the number of matched spikes gets larger.						
Auto Fix	The template is fixed once a set number of spikes have been added. If you have several similar spikes, using Auto Fix after a fairly small number of spikes can stop a template gradually changing shape and becoming the same as another template.						
Track	The template shape tracks the spikes. The contribution of each spike to the template decays as more spikes are added. This is only really useful for slow changes in spike shape and always brings with it the danger than all your shapes will merge together. It also slows down on-line spike classification.						
<i>Spikes for Auto Fix/Track capture modes</i>	This sets the number of spikes for the previous field in Auto Fix and Track modes. In Track mode, the smaller the number, the more rapidly the template shape changes.						
Waveform data	The final group of fields control how the raw waveform data is processed into spikes:						
<i>Waveform interpolation method</i>	Spike waveforms are shifted by fractions of a sampling interval to align them with templates using linear, parabolic or cubic spline interpolation. Linear interpolation is the fastest, cubic spline interpolation is the slowest. The on-line 1401 code always uses linear interpolation for speed reasons. For off-line work, speed of computation does not matter.						
<i>High-pass filter time constant</i>	Use this to remove baseline drift. Set this to a few times the width of the spikes. Do not set the value too low or it will significantly change the spike shapes. If your signal has abrupt baseline changes, you may get better results with the DC offset option.						
<i>Remove the DC offset before template matching</i>	Check this field to subtract the mean level from each spike before matching. This effects template formation and matching, not the saved data. Unless your baseline has sudden DC shifts, it is usually better to use the high-pass filter to follow signals that drift. There is a small time penalty for using DC offset removal on-line.						

Off-line template formation

The Analysis menu **New WaveMark...** command extracts WaveMark data from waveform channels (and also from existing WaveMark data). If you wish to edit WaveMark data, see the *Off-line template editing* section.



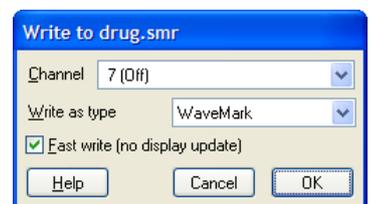
This window is almost identical to the on-line template setup window. There are more menu items and buttons in the control bar, some button labels have changed and there are new controls to set the spike replay rate and you can print and copy template data.

The program scans the data in the channel set by the channel control (wrapping round to the start of the data when it reaches the end if **Circular replay** is checked), and displays triggered and non-triggered events, as if the data were being sampled.

Replay rate This control sets the maximum number of spikes per second to process when running continuously. With **Fast** selected, the spikes are processed as fast as possible; you can also set this by pressing the **F** key (**F** for Fast). With **Real time** selected, spikes are replayed no faster than real time; you can set this with **R** for Real time. You can also select slower rates in Hz; the **S** for Slow key selects a medium-slow rate.

New Channel...

Once you are satisfied with your templates, click the **New Channel...** button. A dialog opens in which you can select the channel to write to and the new data type. The **Fast write** checkbox suppresses display updates whilst producing the new channel, to save time. When you choose **OK**, the selected time range of data is analysed, and any events that cross the trigger levels are written to the new channel in the selected format. Options that change the number of points saved per WaveMark are disabled during analysis. The title of the new channel is set to $nw-c$ where c is the original channel number. The channel comment for the new channel is set to indicate the source of the data.



Time range



This button opens a dialog in which you set the time range of data to process. You can either type in the start and end times, or select values from the drop down list. When you open the window for the first time, the time range is set to the whole file.

Cursor 0

While this window is open, cursor 0 is added to the time view associated with the data file. This cursor marks the position of the current spike. When paused, you can drag this

cursor to a new position. The special cursor does not appear in Cursor windows. The `Ctrl+J` key combination jumps Cursor 0 to the start of the time range.

Circular replay This check box enables circular continuous replay. If you clear this box, continuous replay stops when it reaches the end of the time range set for analysis. The short-cut key for this is `Ctrl+Q`.

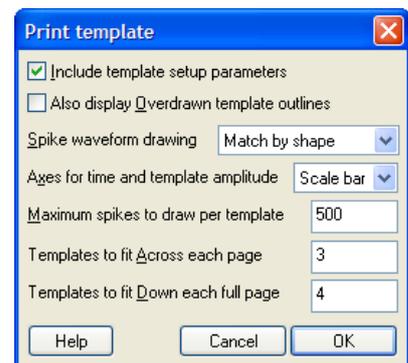


If you depress this button, the associated time window scrolls to keep the current spike centred, if this is possible. The short-cut key for this is `Ctrl+K`.

Marker filter The Analysis menu Marker filter command (short-cut key `Ctrl+F`) is enabled when the source channel holds WaveMark data. It opens the marker filter dialog, or brings it to the front if it is already open. Whenever you change channel to a channel holding WaveMark data with the Marker filter dialog open, the Marker filter dialog automatically changes channel to match.

Print templates The File menu Print... command (short-cut key `Ctrl+P`) opens the Print template dialog. The output format depends on the printer page size (you can set the print margins in the main application File menu Page setup dialog). You can choose how many templates to draw across and down each page and if axes are needed.

You can include the template setup parameters, and also an extra picture showing all of the template outlines overdrawn for easy comparison.

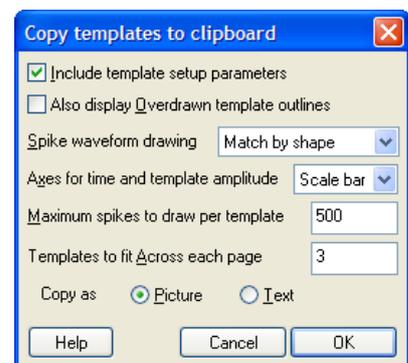


The Spike waveform drawing field chooses how spikes are included in the result. Set it to None for no spikes or to the method for selecting which spikes are drawn with a template. You can choose Match by shape or Match by code (only in the Edit WaveMark dialog). Spikes are drawn in the best-fit position to the template. You can limit the number of spikes to draw in each template; some printers cannot cope with huge numbers of lines in the printed output.

If you have set a time range to analyse, the overdraw spikes are taken from the time range. If you have not set a time range, the spikes are taken from the start of the file.

Copy to clipboard The Edit menu Copy command (short-cut key `Ctrl+C`) opens the Copy templates to clipboard dialog, which is very similar to the Print templates dialog. There are radio buttons to choose between Picture and Text output.

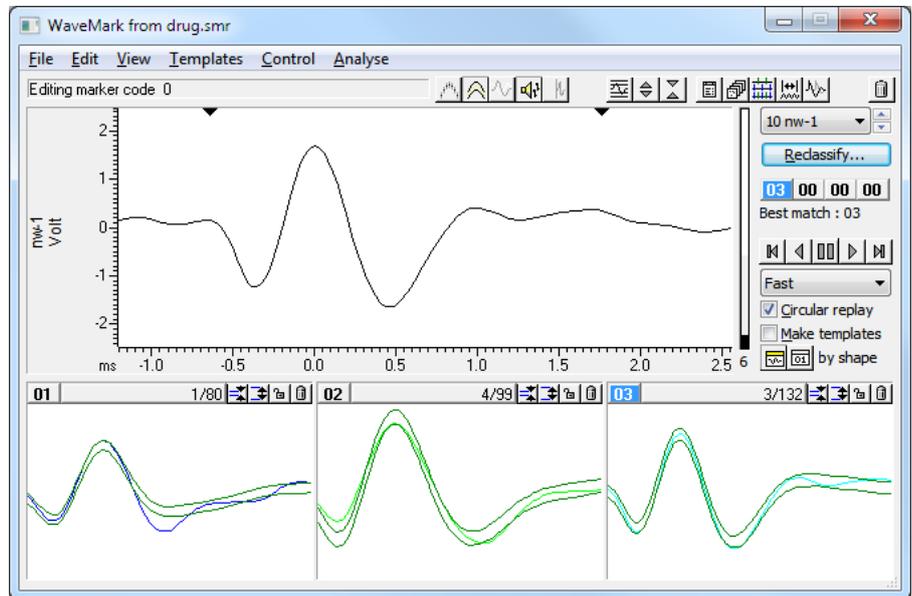
Picture output is in Enhanced Metafile format and can be pasted into documents in most drawing programs and many word processors. Many bitmap-editing programs will accept this format and convert it into a bitmap.



Text output optionally includes the setup parameters. The templates are output in vertical columns separated by Tab characters. There is one column or two columns per template, depending on the template shape display mode (mean templates or upper and lower limits). The first row of template output holds column titles, in quotation marks. The output is designed to paste easily into spreadsheet programs.

Off-line template editing

The Analysis menu Edit WaveMark... command opens a dialog in which you reclassify data in a WaveMark channel. This dialog has its own menu system that includes links to Principal Component Analysis and clustering. If the selected channel has a marker filter set (see the *Analysis menu* chapter), you view and edit only those events that match the filter specification. The window controls are the same as for *Off-line template formation*, except that you cannot change the number of points in the main window and there are no horizontal cursors (use *New WaveMark* if you need trigger levels).



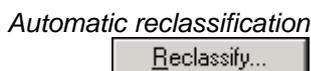
These four codes show the four marker codes of the current spike. Normally only the first code is used for spike classification and the others are 00. If you click on a code, it becomes highlighted and sets the code that is changed by reclassification. If you double-click a code a dialog box opens in which you can edit the code. The **Best match** code is the code of the template that best matches the current spike or 00 if no template matches. If you use this dialog on-line only the leftmost button is enabled.



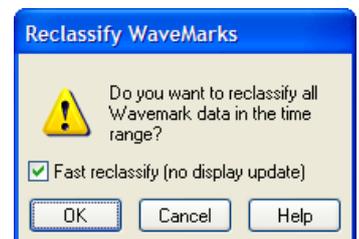
You can build templates by shape or by code. If you build by shape, templates are formed and spikes are sorted based on shapes; prior classification is ignored. If you build by code, the shapes are ignored, and the classification is purely by the code highlighted in the Marker codes. Templates are built automatically when the play controls are set to play forward or backward through the data and when the *Make templates* box is checked.

Manual reclassification

When replay is stopped, you can reclassify a spike by dragging the raw data and dropping it on a template or by double clicking a **Marker code** and editing. You can also reclassify by pressing keys 0 to 9 to give a spike template marker codes 00 though 09. If you use keys 1 through 9 and no template exists with the code, a new template with the code is created based on the current spike.



This option reclassifies all the events (or filtered events if a marker filter is set) on the channel in the time range by matching them to the templates. You can set fast reclassification (no display) or a slower mode where each event is drawn in the data display area as it is reclassified. You can stop the reclassification before all spikes have been scanned, but you cannot undo it.



Tip: clear the **Make templates** check box before reclassifying to stop Spike creating extra templates and modifying existing ones.

Analysis menu commands The Analysis menu contains a range of commands for cluster and Principal Component Analysis, plus commands to manipulate the events that you work with:

Principal Component Analysis This command (short-cut key `Ctrl+A`) initiates PCA analysis of all the spikes in the current time range using the same waveform area as selected for the templates. This is dealt with in the *Clustering* chapter.

Cluster on measurements This command (short-cut key `Ctrl+Shift+A`) opens the Cluster setup dialog in which you select measurements to take from each spike in the time range for cluster analysis. This is covered in the *Clustering* chapter.

Marker filter  The Marker filter command (short-cut key `Ctrl+F`) opens the marker filter dialog, or brings it to the front if it is already open. Whenever you change channel to a channel holding WaveMark data with the Marker filter dialog open, the Marker filter dialog automatically changes channel to match.

Duplicate  This command (short-cut key `Ctrl+D`) creates a duplicate data channel in the time view for each template code in the dialog. It is disabled if the current channel is a duplicate channel. It does the following:

1. All duplicates of the current channel are deleted (with no warning). If you have duplicated the time view, the duplicate channels vanish from all the views. Any processing that depended on the duplicated channels is cancelled.
2. Spike2 counts how many different template codes you have defined. For example, if you have created four templates with codes 01, 02, 02 and 03, this is counted as 3 different codes.
3. For each template code, Spike2 generates a duplicate channel with the marker filter set to display only spikes that match the template code. The marker filter mask used is the same as the code highlighted in the Marker codes (usually the first mask is used). The channel title of each duplicate is set to "title-*nn*", where *title* is the original channel title (shortened to 6 characters if too long) and *nn* is the template code. The new channel is positioned next to the original channel in the time window. It is placed above the original channel unless the Edit menu preferences have reversed the channel order, in which case it is placed below the original.

The assumption is that you have already, or are about to, classify the spikes in the channel based on the templates in the dialog. If you have duplicated the time view, the new duplicate channels exist in all views, but are only displayed in the view that is linked to the Edit WaveMark dialog. The original channel remains unchanged in all the views.

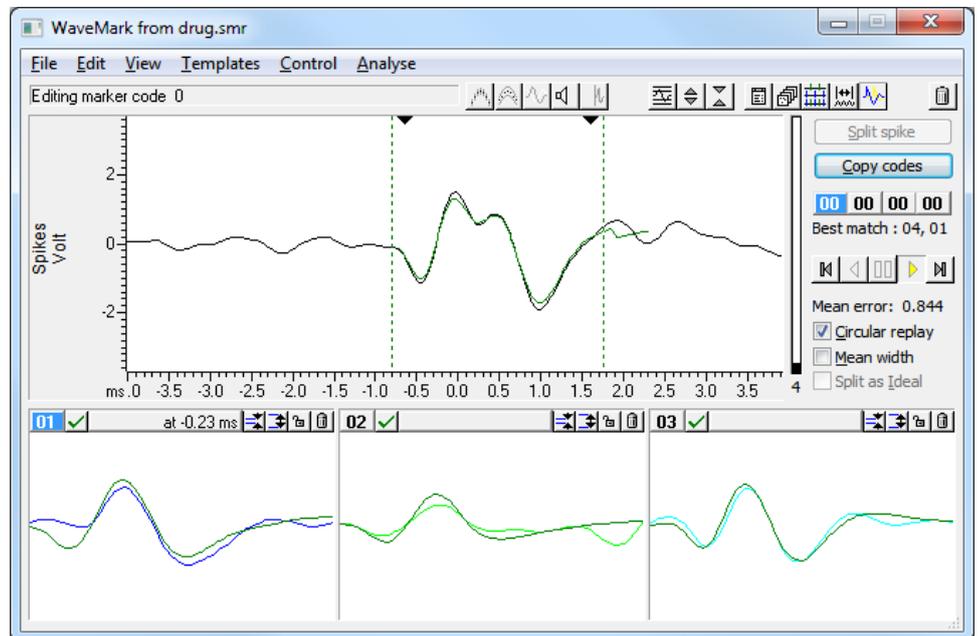
Set Codes This command (short-cut key `Ctrl+E`) opens the Set Marker Codes dialog. You can use this to give all the displayed spikes in a channel the same marker code. For example, if you want to change all the spikes on channel 2 with codes 04, 05 and 09 to have code 03:

1. Open the Marker Filter dialog and set the filter for channel 2 to display only codes 04, 05 and 09.
2. Click the Set Codes button and set the first code to 03 and click the Set button.

Collision Analysis Mode  This command (short-cut key `L`) toggles between normal editing mode and collision analysis mode. Normal editing mode attempts to match each spike to a single template. Collision analysis mode attempts to match the spike to a collision of two templates.

Collision Analysis Mode

If a channel contains more than one class of spikes, and the spikes are independent, it is inevitable that there will be collisions. A simplistic analysis shows that if two spikes have widths w_1 and w_2 and mean firing rates r_1 and r_2 , we would expect $r_1 * r_2 * (w_1 + w_2)$ collisions per second. If both rates are 10 Hz and both widths are 2.5 ms, you would expect 1 collision every 2 seconds. If we can assume that the resulting waveform is the sum of two spike waveforms, and that these waveforms are similar to the established templates, we can search for the template pair whose sum is most similar to the collision.



To enter or leave Collision Analysis Mode, use the Analysis menu command or click on the collision analysis button at the right-hand end of the toolbar. You will notice that some of the controls on the right of the window change, as does the main data display:

Main display The main display shows the current spike overlaid with the best match combination of templates. The dashed vertical lines show the start and end of the template-forming region. The two black triangles mark the start and end of the *Important area*.

Important area You can change the important area by clicking and dragging the black triangles. If the area is smaller than the templates (the usual case), it is always included in the best match. If it is more than twice the size of the templates, the best match area will lie within it. In an intermediate case, it marks the area we would like to match and some or all of the best match will lie in the region.

Template display area The button to the right of the template code excludes the template from the matching process. You might want to do this if a collision is too close to another spike of the same class. You can exclude all but 1 template. Templates used to make a best match have the template code highlighted and list the offset in milliseconds into the main display at which the template starts.

Best match This lists the template code or codes that were used to make the best match trace in the data display. The code for the earlier template is listed first. Templates are allowed to overlap the start and end of the displayed spike when making a match, but at least half the template points must be used in the match.

Copy codes Click this button to copy the *Best match* codes to the current spike. The first code is applied to the currently selected marker code. If there are two codes, the second is applied to the next marker code to the right (wrapping round to the first if necessary).

Mean error This field displays the mean square error per point. The error is scaled by either the template width at each point, or by the mean template width over all the templates.

Mean width If you check this box, the errors are scaled by the mean width of all the templates (including excluded templates). This gives all points the same importance when calculating the best template combination. Think of this as "least squares" fitting. If you do not check the box, the template width is used point by point to scale the error. This gives points with a small template width more importance than points with a larger template width. Think of this as "Chi-squared" fitting.

Split spike This button is enabled if the source channel of the spikes is a memory channel and the current spike can be usefully separated. It replaces the current spike with spikes aligned to the best match templates. How the replacements are calculated depends on the *Split as Ideal* checkbox. If the best match doesn't include the x axis zero, the original spike (less the best match waveform) is also preserved and *Split as Ideal* is assumed to be checked.

Split as Ideal If this is checked (or assumed checked), the replacements waveforms are the template shapes. If this is not checked, the replacements share half the difference between the original data and the best match so that the sum of the two spikes recreates the original where the two created spikes overlap. In both cases, any data that is not available from the original data is set to zero.

Caveats It is important that the templates are centred vertically on zero. If this is not the case, adding templates will create offsets, and the results will not be useful. Likewise, each template should start and end around zero. If they don't, you probably have not selected enough data when forming the templates. If templates start and/or end with a substantial non-zero value, the non-zero ends cause discontinuities in the best match waveform. You must also have sampled fast enough that the change between two consecutive samples is not so large that the matching algorithm fails.

This method will work best in cases where you have a small number of well-defined templates. If you have a large number of indistinct templates, you will likely be able to match just about any shape, but the results will be meaningless. We suggest that you only use this method where there is an obvious collision and that you treat the results with great caution. Look at the firing patterns of the spike classes and check that the resolved collisions fit the patterns. It can be revealing to exclude a template from a collision to see how close the analysis can get with another pair of spikes.

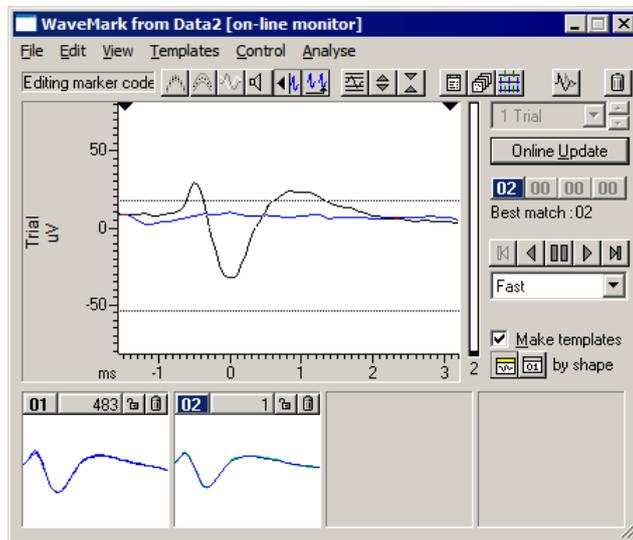
The matching algorithm The best match is found by an exhaustive search of all possible pairs of templates at all alignments against the spike data identified by the *Important area*. The searched area may be further restricted by the rule that at least half the points in each template must overlap with spike data points. Templates are combined by adding the mean template shapes and calculating a width. If *Mean width* is checked, the width is the same at all points and is the average width of all templates. If *Mean width* is not checked, non-overlapped areas of the combination copy the original width and overlapped areas use the RMS (root mean square) of the two widths.

For each pair of templates, there are two variables to consider: the offset of the second template with respect to the first, and the offset of the combination into the data. We find the best value of both variables using integral point shifts, then improve the best value by finding the best fractional position of both variables by interpolation.

The match criterion is the mean squared error per compared point. At each point, the error is the difference between the spike and the template combination divided by a width. The width is either the template width at that point, or is the average template width over all templates. The average width includes excluded templates so that errors are comparable when templates are excluded.

On-line template monitoring

The Analysis menu Edit WaveMark... command can be used during data sampling if there are any WaveMark channels. The window that opens is very similar to the off-line template editing window; the time range and reclassify buttons are hidden, only the first of the Marker code buttons is enabled, the Reclassify... button is renamed to Online update and there is a new button in the control bar.



When this button is down, all spikes are taken from the end of the file, the horizontal cursors that set the trigger levels appear and controls that cannot be used are hidden. In this state, you monitor the spikes as they are classified by the 1401. If there are no spikes, the data display shows recent waveform data so that you can adjust the trigger levels.

With the button up, the window behaves very similarly to the normal off-line template editing except that the Reclassify command is not available. Once sampling finishes the at end of file button vanishes, the Reclassify... button appears and the normal off-line editing behaviour is restored.

Online update

During sampling, the CED 1401 interface classifies spikes based on the templates and parameters that were last loaded to it. If you generate more spike templates, change the existing templates, or edit the template parameters this will make no difference to the 1401 unless you click the Online update button. This deletes all templates stored in the 1401 for the current channel and replaces them with the templates displayed in the window. The trigger levels are updated dynamically; you do not need to click the Online update button to change the trigger levels.

On-line and Off-line templates

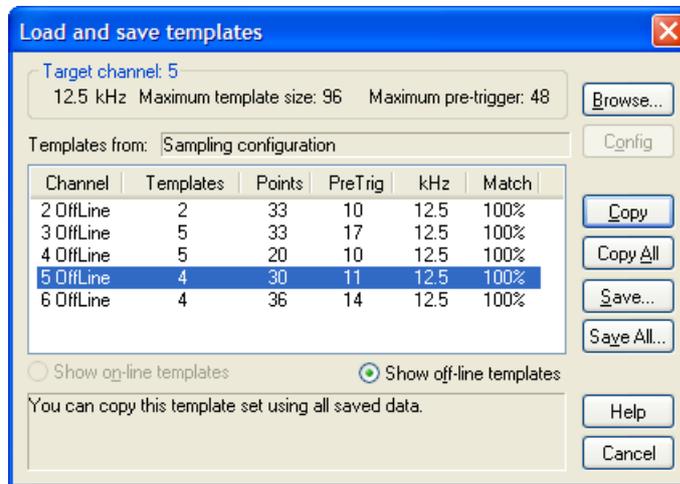
We save two types of template in the sampling configuration and in the S2R resource files associated with a data file: *on-line* templates and *off-line* templates. Spike2 saves on-line templates when you use the On-line template setup dialog and when you click the Online update button while monitoring spikes during sampling. Off-line templates are saved when you change channel or close the spike shape dialog.

When you change channel, Spike2 loads templates from the saved on-line templates if you are sampling and from the saved off-line templates if you are not sampling. This is so that the templates you see during sampling are as close as possible to the templates that the 1401 uses. Off-line templates are loaded from the resource file associated with the data file. If none are found or there is no resource file, off-line templates are loaded from the sampling configuration.

In summary, on-line templates are the last set of templates copied to the 1401 during sampling. Off-line templates are the last set of templates displayed in a template dialog.

Load/Save templates

You can save templates and load templates into the current channel or into all matching channels with the **Load/Save...** button. Spike2 saves templates in the sampling configuration and in the resource files associated with data files. Whenever you work on templates in a data file, for example `data1.smr`, both the associated resource file `data1.s2r` and the sampling configuration are updated to hold the latest set of templates. You can save the sampling configuration to disk in the File menu.



The box at the top of the dialog shows the sample rate and template size limits of the current target channel. This is the channel that is updated by the **Copy** command.

The box in the centre of the dialog lists sets of templates from the Sampling Configuration or loaded from a resource or configuration file. When you open the dialog, Spike2 shows you templates in the Sampling Configuration. The columns are:

Channel Identifies the channel from which the templates originated.

Templates The number of templates stored for this channel. Channels with no templates are not listed. All templates for a channel have the same sampling rate, points and number of pre-trigger points.

Points The number of data points in each template.

PreTrig The number of data points in each template that lie before the peak or trough used as a trigger. A negative value means that the template starts this many points after the trigger point.

kHz The sample rate of the channel from which the templates originated. This need not match the target channel; Spike2 uses cubic spline interpolation to compensate for sample rate differences. Templates generated before Spike2 version 3.15 did not include a sampling rate and 25 kHz is assumed.

Match This column displays how much of each template shape would be used if this set of templates were copied to the target channel.

Browse... and Config Use the **Browse...** button to list templates in resource (`.s2r`) or sampling configuration (`.s2c`) files. The **Config** button displays templates stored in the sampling configuration.

Show on/off-line templates Templates copied to the 1401 during sampling (on-line templates) are saved separately from the templates displayed in the template dialogs (off-line templates); you choose which to list with **Show on/off-line templates**. The **Copy All** and **Save All** commands operate on the list selected by this option.

Copy The **Copy** button copies the selected set of templates into the target channel and closes the dialog. You can copy a channel as long as **Match** is at least 10%. This overwrites any templates that already exist for the target channel.

- Copy All** For each channel in the channel control, Spike2 searches the list of templates for a matching channel number. If templates exist for the channel and they match the channel settings, they are copied from the list and become available for use.
- Save** This saves the selected set of templates in the list to a resource file (.s2r file extension). You are prompted to supply the name of a new file or to select an existing resource file. If the resource file already exists, any templates for the channel are replaced. If the resource file does not exist, the templates are written to a new file.
- Save All** This saves all the templates in the list to a resource file (.s2r file extension). You are prompted to supply the name of a new file or to select an existing resource file. If the resource file already exists, any existing templates for the channels in the list are replaced. If the resource file does not exist, the templates are written to a new file.
- Cancel** This button closes the dialog.

Template storage

You use templates to classify WaveMark data, but they are not stored in data files. Instead, Spike2 stores templates in the following places:

- In the sampling configuration. You can always load the last set of templates you used on-line or displayed off-line in the current session from here. You must save the configuration to make templates persist after you close Spike2.
- In saved configuration files with the file extension .s2c. For example `Last.s2c` automatically saves the last configuration used for sampling. However, changes made in the template dialogs after sampling ends are not saved automatically.
- In resource files associated with data files. Each time you use a template dialog, any templates you create or change are saved in a resource file with the same name as the data file, but with the file extension .s2r. During sampling, there is no associated resource file. In this case, when you save the data file, Spike2 creates a resource file and copies the templates in the sampling configuration to it.
- In resource files created from the **Load and Save templates** dialog.

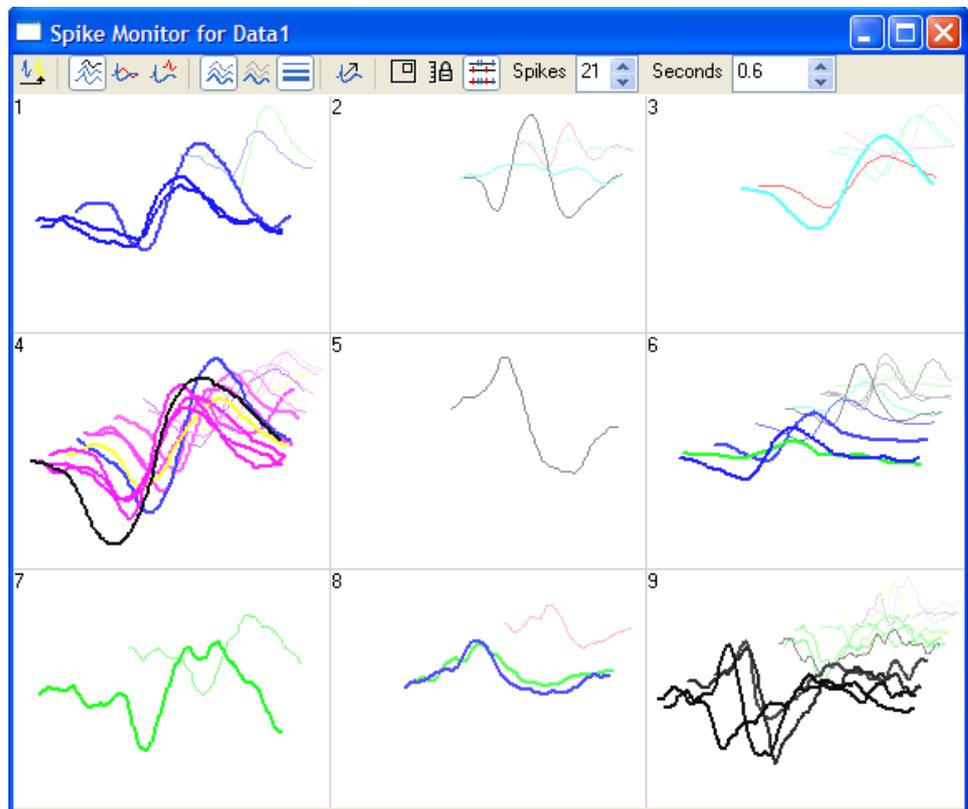
When you open a new data file for sampling, the template dialogs load up the last set of on-line templates from the sampling configuration. These on-line templates are updated each time you click the **Online update** button in the **Edit WaveMark** dialog. The off-line templates are updated to match the templates displayed in the template dialog. When you stop sampling and save the file, the templates are written to both the resource file associated with the data file and to the sampling configuration.

When you work on a data file from disk and open a template dialog, the off-line templates are loaded from the resource file associated with the data file. If none are found, off-line templates are loaded from the sampling configuration. Any changes you make are saved to both the resource file and to the sampling configuration as off-line templates.

By saving templates to both the resource file and to the sampling configuration, Spike2 makes it easy for you to work on a sequence of data files. For example, you might break a sampling session into several data files. On-line templates created for the first data file are automatically used for the second and subsequent files.

Spike Monitor

The Spike Monitor window can be opened with the View menu Spike Monitor command or from the script language with `SMOpen()`. The main use of the window is to give an overview of the current spike activity on all channels during sampling. You can drag and resize the window to position it; the grid of cells organises itself automatically. Each cell displays the channel number at the top left and spikes that lie in a user-defined time range back from the current time. The *WaveMark background* field of the *Application colours* sets the cell background colour.



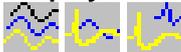
The window is linked to the Edit WaveMark dialog. Double-click a cell to open the Edit WaveMark dialog. Single-click a cell to select the corresponding channel in the Edit WaveMark dialog. The display is controlled by the toolbar at the top of the window. The toolbar commands are:

At end



This button is enabled for on-line use and during rerun off-line. With the button down, spike data comes from the end of the file. With the button up, the data comes from the cursor 0 position in the associated time view. If you start to rerun a data file, this window automatically switches to *At end* mode, and cancels this mode when the rerun ends.

Display mode



There are three display modes: *3D*, *2D* and *2D separate*. Spikes are positioned in the cell using two rectangles: a front rectangle that is anchored to the bottom of the cell and a back rectangle that is anchored to the top of the cell. You can display the rectangles and change their positions with the Rectangle command, described below. Normally the front rectangle is larger than the back rectangle. The display modes are:

3D Each spike is drawn in a rectangle that is positioned between the front and back rectangles depending on how old the spike is relative to the current time. New spikes appear in the front rectangle and then move to the back rectangle as time passes.

2D All spikes are drawn in the front rectangle; the back rectangle is not used.

2D separate The last spike is drawn in the front rectangle, the rest are drawn in the back rectangle.

Colour fade



With this button down, the spikes colours fade into the background colour with time. A new spike is drawn in its normal colour. A spike that is half the user-defined time range old would draw in a colour that is half way between the normal colour and the background colour.

Colour front only



With this button down, only the newest spike is drawn in its normal colour. All other spikes appear in a grey that contrasts with the background enough to be visible. This is particularly useful in the 2D drawing mode.

Vary line thickness



Normally, all spikes are drawn with the same line thickness as data in a time view as set in the **Edit** menu **Preferences**. With this button down, the lines become thinner with time, with new spikes being drawn with a thick line. The line thickness for a new spike depends on the height of the cell. Use this option with caution; a thick line takes longer to draw than a thin one.

Timed (smooth) updates



In *At end* mode, cells update when new spikes enter or leave the time range, which can lead to a jerky display. With this button down, cells update continuously, which can look much better, particularly in the 3D display mode. However, this option uses more system time especially if *Vary line thickness* is enabled. The continuous update rate is decreased when the Spike Monitor window is not the active window to make more time available to other windows.

Show rectangles



Normally, the front and back rectangles that are used to position spikes are hidden. With this button down, the rectangles are drawn and can be moved by clicking and dragging with the mouse. The spikes are still displayed, but are drawn in grey so that the rectangles are clearly visible. To move a rectangle sideways, click in it and drag it. To size a rectangle, click in it to display two drag handles and then click on a drag handle to resize.

Lock y axes



With this button up, the spikes in each cell self-scale. Each cell tracks the maximum amplitude it has seen. When a new spike that exceeds the current maximum amplitude is seen, it updates the maximum amplitude. This new maximum is held for two seconds, then it decays with time. With this button down, the current y axis scale is held.

Duplicate channels



With this button up, duplicate channels are not displayed and any marker filters set for channels are ignored. With this button down, duplicate channels are displayed and only the spikes that meet the marker filter settings for each channel are displayed.

Spikes

You can set the maximum number of spikes to display in each cell, in the range 1 to 40. The time taken to display the data depends on the number of spikes, especially if you enable *Vary line thickness*. You should set this to the minimum number that is useful.

Seconds

This field sets the time back from the current time range over which to display spikes. It is normal to set this to a small number of seconds, but you may need to set a very short period if you are trying to visualise burst characteristics in the 3D display mode.

Getting started with spike shapes and templates

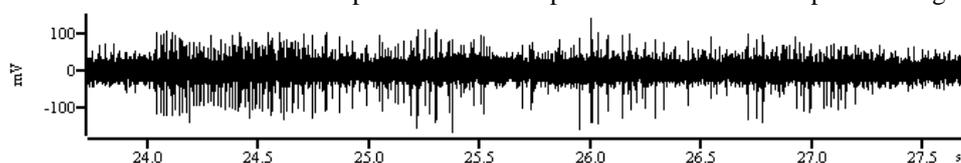
Sorting spikes is as much an art as a science; there is no substitute for the skill and experience of the investigator. Spike2 provides you with a toolbox of routines that will help you to discriminate various shapes, but there is no guarantee that two spikes are from the same source just because they have the same shape (or because any other measured parameter or set of parameters are the same).

Quality of the original signal

Probably the most important contributor to successful spike sorting is the quality of the original data. If you can improve the signal to noise ratio, or adjust the electrode position to increase the amplitude of your target spikes relative to the background, the time spent will be rewarded by more certain classifications in less time.

The waveform sample rate used is also very important. For most spikes lasting 1 to 2 ms, a sample rate of 20 to 25 kHz is about right (some workers like higher rates than this). It is important that the data is frequency band limited to half the sample rate before it is sampled. For example, when sampling at 25 kHz there must be no frequency components above 12.5 kHz. If such components are present, they are aliased to lower frequencies and appear as noise that is very difficult to remove.

To check the quality of your signal it is a good idea to sample a section of data as a waveform first. Here is an example waveform sampled at 25 kHz from a tape recording.

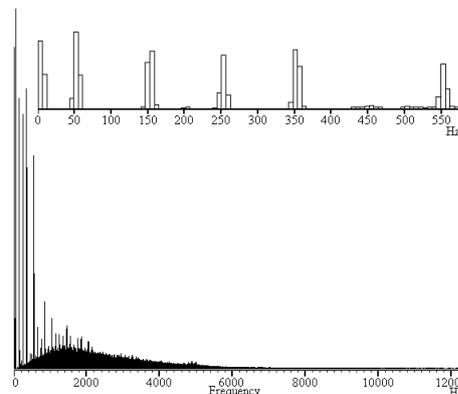


Sampling configuration for this data

Open the sampling configuration Channels tab and click Reset. Click New Channel and select a Waveform channel. Set Channel to 1 and 1401 Port to 0. Set the Ideal adc sampling rate to 25000, and set Units mV = Input in Volts x 1000 + 0 and click OK. Now click the Resolution tab and set Optimise to Full. Finally click the Mode tab and select Continuous. You can now sample.

The example waveform amplitude is around ± 100 mV, which is a bit small. The 1401 family usually expect a full-scale input range of ± 5 Volts, and our spikes are using only 1/50th of the available range. The size of each digitisation step is about 2.5 mV for the micro1401 and 1401*plus* and 0.15 mV for the Power1401 or Micro1401. Although the system will work fine with signals of this amplitude, if your rig has more gain available, it is much better to make the signals around ± 1000 mV in amplitude. This gives you plenty of headroom for large spikes and a good signal resolution.

Check the input for unexpected frequencies with the Analysis menu: New Result View: Power Spectrum command with a block size of 4096. Here we found narrow peaks at low frequencies and almost nothing above 6 kHz. Apart from a peak at 0 Hz caused by a small DC offset, the remaining peaks are all at odd multiples of 50 Hz. This data was recorded in England where the mains frequency is 50 Hz. Anything that reduces mains pickup (better shielding, removal of earth loops, cleaning up of connections) will make spike sorting easier.



Sharp power spectrum peaks often indicate signal contamination. If the peaks change frequency when you change the sample rate, they are caused by aliasing of frequencies above half the sampling rate and they may be removable with an external low-pass filter.

Setting up the WaveMark channel

Assuming that we have adjusted the rig to get the best signal to noise ratio and spike amplitude, we must now set up a suitable sampling configuration to capture our spikes. We could sample the data as a waveform, but unless you really need all the data in between your spikes, this approach is wasteful of disk space. For example, a single waveform channel sampled at 25 kHz consumes 50 kB of disk space per second, or 3 MB per minute, or 180 MB per hour.

An alternative is to sample WaveMark data. In this case, each time the input signal crosses a positive or negative threshold we search for a peak or trough in the data and save the time and a small fragment of waveform. If we save 32 data points per spike, each spike costs us 72 bytes of data. As long as you set the threshold levels so that you get less than 700 spikes per second, this method uses less storage space. At a mean spike rate of 20 spikes per second, you would consume 86 kB per minute, or 5 MB per hour.

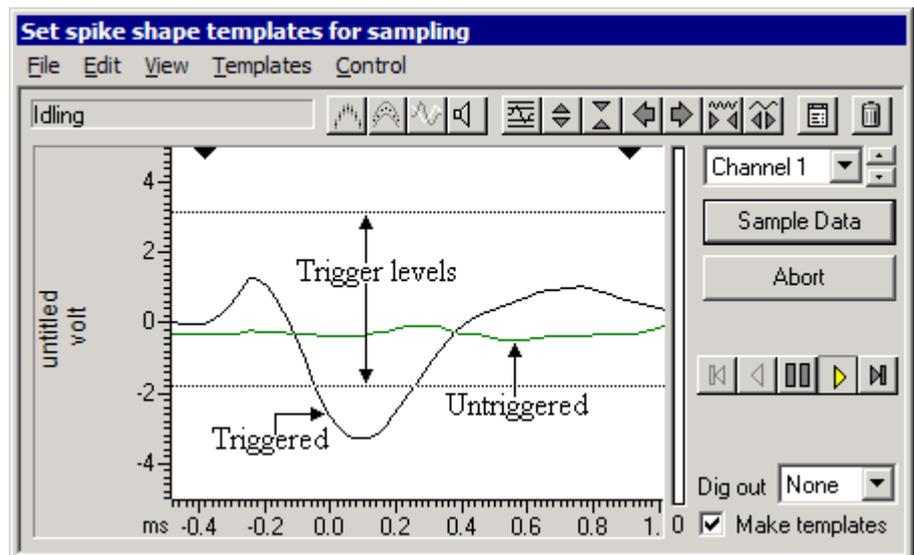
There is one further advantage to WaveMark data. As the data is captured, Spike2 can match it to templates and code each spike accordingly. This gives you further on-line analysis capability as you can look at the responses of individual spike types on-line. The on-line sorting is not the final word as you can review and resort the data off-line.

Suggested configuration for WaveMark data

Open the sampling configuration Channels tab and click the Reset button. Then click New Channel and select a WaveMark channel. Channel should be 1 and 1401 Port should be 0. Set the Maximum sustained event rate to 30, and set Units $mV = \text{Input in Volts} \times 1000 + 0$. Set Points to 32 and Pre-trigger to 10 and WaveMark sample rate to 25000 and click OK. Now click the Resolution tab and set Optimise to Full. Finally click the Mode table and select Continuous. You can now sample.



Spike2 knows that you have a WaveMark channel in the sampling configuration, so when you click the Sample now! Button in the toolbar, it opens up a window for you to set trigger levels and adjust the number of sample points and to build templates.



The picture shows the window after adjusting it to have the minimum size and setting up the trigger levels. When you open a window for the first time you are likely to have a more or less flat line across the centre of the display area and the trigger levels will not be in useful positions.

There are a lot of buttons and controls in this window. If you move the mouse pointer over one and leave it for a second or so, a line of text will “pop-up” with a short explanation. Alternatively, you can use the F1 key to activate the Help system.

Home key or 

This button is one of the most useful when you are setting up for spike shape capture. Each time you click the button, Spike2 will make a guess at reasonable trigger levels (and limit levels for the New WaveMark command with limits enabled) and adjust the display gain so that you can see the data. Once a trigger level is crossed a second trigger cannot occur until the waveform falls below half the trigger level.

It is likely that you will only wish to trigger in one direction, so move the cursor that you do not want to use to the edge of the window. To change a level, move the mouse pointer over the level, hold down the mouse button and drag to the new position.

It is possible to use both trigger levels, but we strongly suggest that you do not to avoid the sideways shift you will get if a spike triggers on the opposite side from the one you intended. If your spikes are bi-phasic, that is they have a rising peak and a falling peak, it is best to trigger in the direction that has the clearer peak. If your spikes are tri-phasic, as in the example above, it is best to trigger in the direction that has only one peak to avoid problems caused by false triggers on the second peak. If both directions are equally clear, choose the direction with the narrower peak (usually the first one).

Up and Down keys or 

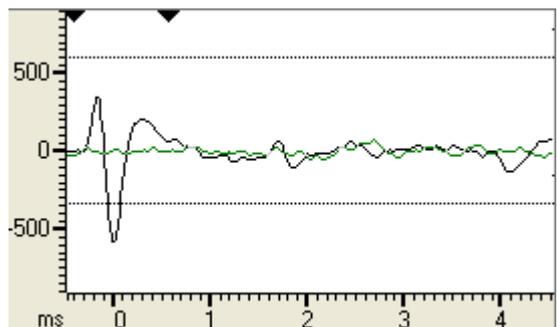
Once you have got more or less the desired display gain you can use these two buttons to change the display scaling manually. You may also want to reduce the display gain so you can drag one of the trigger levels off the screen to reduce visual clutter. As an alternative to using this button, move the mouse pointer over the y axis and click and drag to scale the data.

Left and Right keys or 

These buttons change the number of pre-trigger points. The trigger point is the first signal peak or trough after it crosses the trigger level. This point is labelled 0 on the horizontal axis. You cannot change the number of pre-trigger points after you start to sample. You can also scroll the data by clicking and dragging the tick marks of the x axis.

PgUp and PgDn keys 

These buttons increase and decrease the data points that are displayed and written to disk for each spike. The number of pre-trigger points is preserved if this is possible, so you should set the pre-trigger points first. You can scale the x axis around the zero point by clicking and dragging the x axis numbers.



You need to capture sufficient points to enable accurate spike identification. You also want to capture as few points as possible to minimise the data storage. In this example we have captured about 5 ms of data, which is at least 3 ms more than we need. As well as leading to huge data files, sampling too many points also increases the chance of the data for one spike containing data for a second or even a third spike. Further, Spike2 will not trigger for the next spike until it has finished sampling the previous one. You cannot change the number of points once you start sampling.



The two black triangles mark the start and end of the region of the spike that is passed to the template matching system. Click on a triangle and drag to change the region. When you drag the markers, vertical cursors drop down to identify the selected area.



This is the play control. The buttons from left to right are: step backwards, run backwards, pause, run forwards and step forwards. Step and run backwards are disabled during setup. You can also use keyboard short cuts for these buttons: **b** or **B** steps backwards, **N** runs backwards, **v** or **V** pauses, **n** runs forwards and **m** or **M** steps forward.

Forming templates

By now you should be able to set up the system to capture spike shapes. The next step is to form templates from these spikes. The first step is to check that the template parameters are set to sensible values.

Ctrl+Enter or 

Click this button to open the template parameters. To get started, set these values:

Suggested template parameters

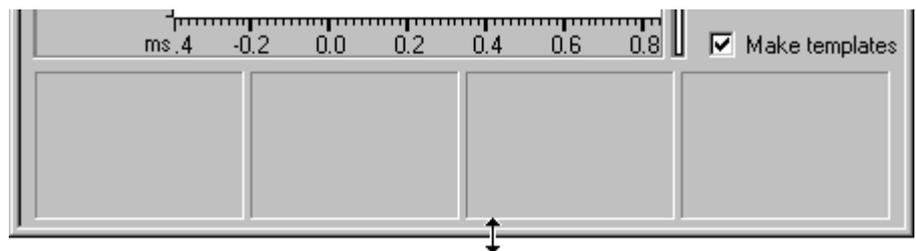
In the **New template** section, set **Number of similar spikes for a new template** to 8, **New template width as a percentage of amplitude** to 30, **No template for spikes rarer than 1 in** to 50. In the **Matching a spike to a template** section, set **Maximum percent amplitude change for a match** to 0, **Minimum percentage of points in template** to 60 and check the box for **Use minimum percent only when building templates**. In the **Template maintenance** section, set **Template modification mode** to **Add All**. In the **Waveform data** section, set the **Waveform interpolation method** to **Linear**, **High-pass filter time constant** to 20 ms and leave the **Remove DC offset before template match** as unchecked. Finally click **OK**.

As you get more familiar with template matching you will probably change these values. However, this set is a useful starting position. We have turned off the amplitude scaling as this causes a large increase in on-line computation and so should not be used unless you really need it. Amplitude scaling comes into its own when your spikes change amplitude while preserving shape, for example during a burst.

Del key or 

Spike templates are displayed below the data display area. If you cannot see the template area drag the bottom edge of the window downwards with the mouse. The template area may already contain templates. If it does, you can clear them all by clicking the bin button at the top right of the window. You can remove individual templates by clicking the small bin button in the template window.

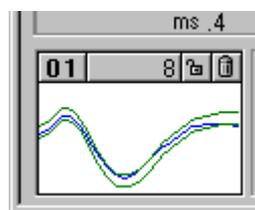
Expanding the window to show the template area



The template windows have two sizes: small (in the picture) and large. The large size is slightly more than twice the width and height of the small image. You change from small to large by double clicking in the template area. Spike2 remembers the last window size you used for both small and large template displays. You can also expand the window sideways by dragging the right-hand edge of the window, or you can grab one of the corners to expand in both directions. As a short-cut, you can maximise and minimise the size of the window by double-clicking in the title bar.

Make templates

To build templates, make sure that the **Make templates** box is checked, the play control is set to run forwards and that spikes are crossing the trigger level you set in the data display area. If all is well, when 8 similar spikes have been detected (8 is the number you set in the template parameters) a new template should appear in the template area.



The four rectangles at the top of the template hold, from left to right, the code associated with the template (**01**), the number of spikes in the template (8), a button that indicates the locked state of the template and a button that can be used to delete the template. The vertical size of the displayed template is set by the vertical scale of the data display area.

As spikes are processed, the vertical bar on the right of the data display area indicates the number of templates that are being considered in grey and the number of confirmed templates in black. The first 20 confirmed templates appear in the template area. The count of confirmed templates is also displayed below the bar.



The horizontal size of the template is set by the two black triangles in the data display area. For good spike sorting, it is important to select the region of the spike that shows the most variation between different classes of spikes. Do not include baseline areas before or after the spike. Spikes are matched to templates based on the errors between the mean template and each spike. Spikes are excluded from templates based on individual points falling outside the template boundaries.



You can control the appearance of the templates with these three latching buttons. If you click them once they become active, click them again and they become inactive. The leftmost button controls overdrawing of spikes in the template. With the button down, all spikes that match the template are over-drawn in the template. With the button up, only the last matching spike is drawn.

The centre button chooses between a display of the template boundary with the button down, and a display of the mean template with the button up. With the right-hand button up, spikes are displayed only in the template that they match. With the button down, spikes are also displayed (in grey) in all non-matching templates.

Spike code and template colour

Each template has a code in the range 1 to 255. When the templates are used for spike sorting, the code is used to label matching spikes. Code 0 is not used by templates, but spikes can be given a zero code, meaning that they are not coded.

The code is represented by two hexadecimal digits; 1 is represented by 01, 2 by 02 and so on up to 9 which is 09. In hexadecimal (base 16 numbers), the numbers ten to fifteen are represented by the letters A to F, so ten decimal is 0A, eleven is 0B through to fifteen which is 0F. Sixteen decimal is 10 in hexadecimal, seventeen is 11 and so on.

There is a further complication with the codes. We also allow the use of single character codes in place of the hexadecimal codes 20 to 7E. This came about because the keyboard marker channel uses the same coding system and the ASCII codes for the Roman printing characters lie in the range hexadecimal 20 to 7E. Although Spike2 can manage up to a hundred templates internally, you can display only the first twenty, so codes above 14 (decimal 20) are rarely used and single character codes are rarely a problem.

The code given to a template also determines the template drawing colour. In the View menu **Change Colours...** dialog you can assign colours for template codes 01 to 08. It is a good idea to set colours that contrast strongly with the background colour you have set for the time view as this colour also sets the background colour for the template windows. If you use higher numbered templates the colours repeat, so the colour for codes 09 to 10, 11 to 18, 19 to 1F and so on are the same as for 01 to 08.

New templates take the lowest available code, but you can change the template code by double clicking on it. This opens a dialog box in which you can edit the code. There is no restriction on the code you give, other than it must be in the range 01 to FF. In particular, you can have more than one template with the same code.

If you enlarge the window on the right, you will discover two more buttons. **ReNUMBER** changes template codes to remove gaps caused by deleting or editing codes. **Order by Code** sorts the templates into ascending code order.

Automatic template creation

With the play control set to run and the **Make templates** box checked, Spike2 creates templates automatically by looking for similar spikes and following the rules set in the template parameters dialog. Each new spike is compared first with the existing confirmed templates (displayed in the template area). If it doesn't match any of them it is compared with the provisional templates. If it matches nothing, it forms a new provisional template.

When a spike matches a template it modifies the mean template shape and width unless the template is locked. If the spike matches more than one template it is added to the one with the smallest error between the spike and the mean template shape.

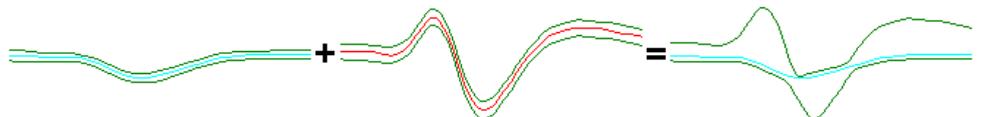
When a provisional template gets a new spike and this causes the template spike count to reach the **Number of similar spikes for a new template** set in the template parameters, it is ready to become a confirmed template. If it matches a confirmed templates it is merged with the confirmed template. Otherwise, a new confirmed template is created.

To avoid the system being swamped with provisional templates, there is a mechanism that makes them decay away. This is set by the **No template for spikes rarer than 1** in field of the template parameters dialog. If this number of spikes are tested and a provisional template does not get one new spike added to it, the template spike count is reduced by one. If the spike count becomes 0, the provisional template is deleted.

Manual template creation and merging

Instead of letting the system create templates, you can build them yourself. Use the play control to step forwards one spike at a time until you find one that you want to turn into a template. Click the mouse in the middle of the display area and drag the spike to an unused window in the template area and release. If you drag a spike shape to an existing template, the spike is added to it.

If you decide that two templates are very similar, you can combine them by dragging one of the templates and dropping it on another. Spike2 will align the dragged template with the target and then merge the two shapes by combining the template limit.

Merging templates

This example shows the effect of merging two very dissimilar templates. You can see that the extremes of the template boundaries of the original templates are preserved in the result. Normally, you would merge two similar templates. You can test for template similarity by dragging a template over another without releasing.

Number of templates

You can create (and Spike2 will remember) up to 20 templates per channel. However, if you are preparing for on-line sampling, only the first 8 templates in the template area will be passed to the 1401 for use on-line. You can, of course, resort the data off-line using all 20 templates, and it is possible to sort spikes into many more classes than this by using the Analysis menu **Marker Filter** and **Edit WaveMark** commands.

Locking templates 

Each time a spike matches a template, and the **Make templates** box is checked, the spike is added into the template unless the template is locked. The lock button in each template can be used to lock and unlock templates manually. You can also make the templates lock automatically after a set number of spikes have been accumulated by setting the **Template modification mode** in the template parameters dialog to **AutoFix**.

If you don't lock the templates, they change as more spikes are added. This means that the shape you end up with could be quite different from that with which you started. For example, if you create a template manually from a specific spike, you may want it to keep the exact shape so you would lock it.

Amplitude variation If you find that the system keeps generating multiple templates for the same spike because the spike amplitude changes, there are several things to try. The obvious solution is to set the **Maximum percent amplitude change for a match** to a non-zero value. For example, if you set this to 30%, each spike can be increased or decreased in size by up to 30% before attempting to match it to the templates. Spike2 does this by computing the area between the spike and the zero baseline and changing the spike amplitude so that this area is the same as the area between the template and the baseline.

Beware that this is computationally expensive, which does not matter for off-line sorting, but it can be important on-line, especially if you don't have a Power1401. The more channels of spikes and the more templates you are matching against, the larger the time penalties you will impose. If you find that turning this option on leads to spikes not being classified when the spike rate is high, then you may need to consider other options.

Another method to cope with amplitude variation is to make the templates wider. You can do this by merging templates, or by increasing the **New template width as a percentage of amplitude** in the template parameters. However, wider templates leads to less discrimination between spike shapes.

Closely allied to making the templates wider is to reduce the **Minimum percentage of points in template** field. However, this also leads to less discrimination between spike shapes.

You might also consider checking the **Use minimum percent only when building templates** box. If you do this, when you run on-line and when reclassifying spikes, spikes are matched to the template that has the smallest error between the shape of the spike and the shape of the template, the template limits are ignored. This option is most useful when you know that a whole bunch of spikes are either type A or B but the data is very noisy, so you chose a good spike of each class to make two templates, then you check the box and sort on the basis of minimum error.

If you do not need to separate many shapes on-line, you can simply accept that it takes 2 or even 3 templates to cover the full range of amplitudes and give the 2 or 3 templates the same code (double click the code and edit it).

Using sound 

If your computer has a sound card, you can play a sound for each spike. Spikes that match a template play a sound with a pitch that rises as the code for the template increases. Spikes that don't match will play a lower pitch. How successful this is depends on the capabilities of your system! It works well on my computer when running under Window NT, 2000 or XP, but on the same computer running Windows 95 or 98 the result is not so good. Try it and see. Ideally, you should be able to distinguish different patterns for different spikes.

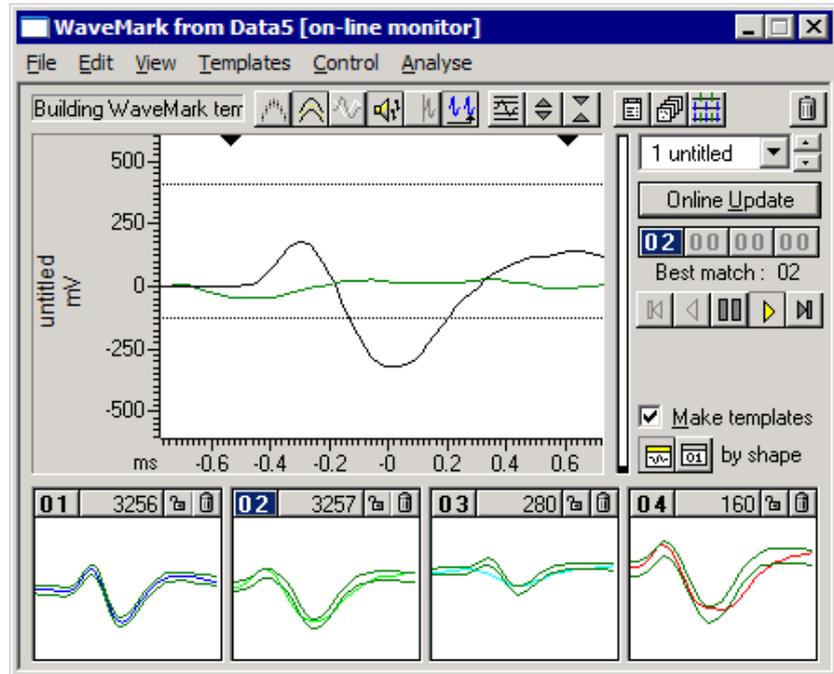
Sampling data

Although we have not discussed the details of every possible control you could have used, you should now have enough information to be able to set up your system with some templates and sample data. Click the **Sample Data** button and the dialog box will close and a new data file window opens. If you followed our sampling configuration settings, the window will contain a channel of WaveMark data and the Keyboard marker channel.

When you close the template setup dialog, you fix the number of points that will be captured for each spike, and the number of pre-trigger points.

Monitoring spikes

You can adjust the trigger levels and even change the templates during sampling with the Analysis menu Edit WaveMark command. If you use this a window that is very similar to the template setup window opens:



The displays area shows the last spike that crossed the trigger level for the channel, and when there are no spikes, it shows the latest data. You can adjust the trigger levels by dragging them, and *the level change is passed to the 1401 immediately*. If you make any other change, this has no effect on the 1401 unless you click Online update.

During sampling, the 1401 holds a set of templates that it uses to classify spikes. There is a separate set of templates held by Spike2 that start out as identical to the ones in the 1401. However, you can change them by adding more spikes, deleting templates or creating new ones. The templates in the 1401 do not change unless you click Online update to copy the latest templates to the 1401.



The leftmost of these 4 buttons shows the marker code assigned by the 1401; the other three buttons are not used on-line and are disabled. The **Best match** field displays the marker code of the template in the template area that is the best match to the spike in the display area. These will usually be the same. They can differ if you have edited the templates, or if **Make templates** is checked so that the template area is changing.



If you enable sound, the tones you hear depend on how the spikes match templates in the template area; the tones do not depend on the classification done by the 1401. This is to allow you to modify templates in preparation for Online update.



This is a latching button that is normally down for on-line use. Click the button to change the state between down and up. With this button down, the displayed spikes are always taken from the end of the data file, that is from the spikes that have just been sampled. If the spike rate is too high to show all the spikes, spikes are skipped.

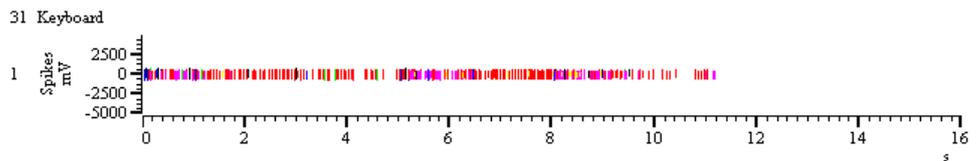
With the button up, a new cursor appears in the time view to mark the point from which spikes are taken, all the play control buttons become enabled and the window behaves almost identically to the off-line Edit WaveMark window which is described later. The marker code buttons still only allow the first marker code to be used; you can use all four codes when off-line.

Monitoring multiple channels

If you have multiple channels of WaveMark data, you can monitor them all by opening the Spike Monitor window (use the View menu Spike Monitor command). This view allows you to check all the WaveMark channels at a glance. This window also has links to the Edit WaveMark dialog. If you click on a channel in the Spike Monitor window, this selects that channel in the Edit WaveMark dialog, ready for adjustment.

Drawing modes for spikes

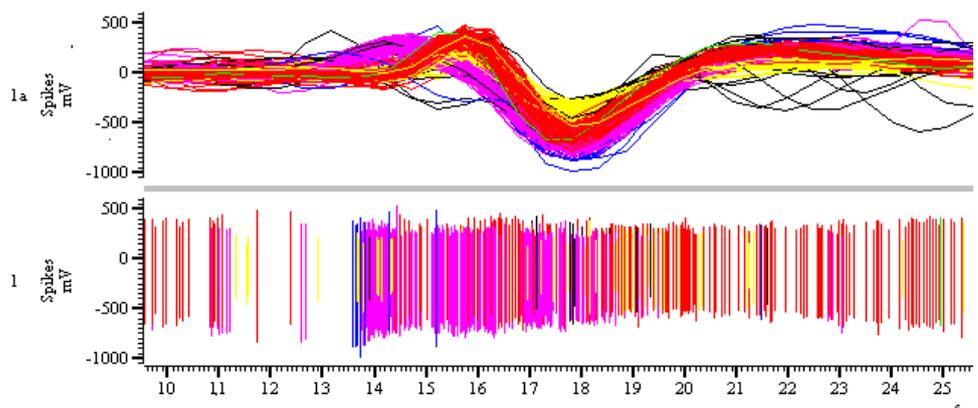
If you use the View menu Channel Draw Mode... command and select the WaveMark channel, you will find that you have a very wide choice of display modes. To start with, display the channel as Waveform. Arrange the x axis to show 10 to 20 seconds worth of data and click the Start button in the sampling control bar or use the Start Sampling command in the Sample menu. You should see something like the following:



If no spikes appear, use the Analysis menu Marker Filter... command and select each of the four layer in turn and click All, then click OK. We will discuss this command in more detail later.

The spikes are colour coded for the template that they match, or are drawn in black if they don't match any template. These spikes are rather small, so the first thing to do is to double click the y axis and optimise the display. You can display the codes for the spikes as well as the waveforms by changing the drawing mode to WaveMark, but this is not very useful when you have a lot of spikes in the window as the codes will overwrite each other. Drawing in WaveMark mode on-line is much slower than drawing as Waveform, so we do not recommend this mode for on-line use.

Another useful mode is Overdraw WM (overdraw WaveMark). This gives an immediate overview of the spike sorting. To see this mode in action, use the Analysis menu Duplicate Channels command to duplicate the WaveMark channel. Then use the View menu Channel Draw Mode... command to set the drawing mode for the duplicated channel to Overdraw WM. You may need to adjust the size of the data window to generate a useful image. In the next example we have hidden the keyboard channel.

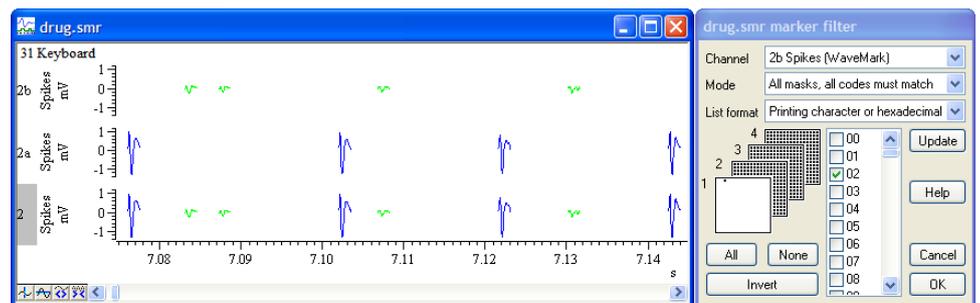


In Overdraw WM mode, the full width of the window is used for each spike. There is a grey bar between the overdraw area and the rest of the window to show that the x axis does not apply. When the window scrolls, new spikes are added to the overdraw area, but old spikes are not removed as this would force the entire window to redraw, which could take a long time. If you want to force the overdraw area to redraw, click or drag the thumb in the scroll bar at the bottom of the window.

You can locate a particular spike in Overdraw WM mode. Move the mouse pointer so that it is over the desired waveform at a place that is clear of all other waveforms, right-click and select Find with cursor 0 from the pop-up context menu. This moves cursor 0 to the position of the event. If the Edit WaveMark dialog is open, it will also move to the position as cursor 0 controls where the dialog collects data.

Using the Marker Filter

Although it is nice to see all the spikes you have collected, for most purposes, you are interested in the behaviour of one unit, or the relationship between one unit and another. If your templating and spike sorting has gone well, you will have isolated each unit as spikes with a particular code (or possibly more than one code). You now need a way to treat each unit as a separate channel of data.



The hard way

In the example shown, channel 2 has two classes of spike, sorted into codes 1 and 2. You could generate channels 2a and 2b by duplicating channel 2 twice (right click on the channel and select Duplicate from the context menu). Then right click on channel 2a and select the Marker Filter command, clear the check box next to 01 and click the Invert button. This clears all check boxes except the one next to 01. Then click Update and channel 2a will display only spikes with code 01. Next, select channel 2b in the channel list and repeated the procedure for code 02.

The easy way

However, there is a much easier way to do this. The Edit WaveMark dialog has a button that generates a duplicate channel with the marker filter set to display each template code. Alternatively, you can use the Edit WaveMark dialog Analysis menu Duplicate command (short-cut `Ctrl+D`).

The result of this is that we have two new channels, each with spikes from a single unit. If your sorting resulted in one unit having more than one code, then you should select all the codes for that unit in the Marker Filter dialog; there is no need to sort spikes so that one template matches every single spike from the unit.

If you decide that spikes with several different codes should all have the same code, you can use the new Set Marker Codes command in the Analysis menu, or right click on a channel and select this command from the context menu. Set the marker filter for the channel to display the codes that you wish to amalgamate then use the Set Marker Codes dialog to change all the spikes with these codes to the same code.

Creating on-line templates with cluster analysis

Cluster analysis is normally used off-line. This is because extracting principal components and automatic clustering can take a noticeable time if you are working with a large number of events; also the clustering analyses are not available from the on-line template setup dialog. However, you can use clustering to create on-line templates. There are two ways to do this:

Off-line clustering Sample as normal and set the spike trigger levels in the on-line template setup dialog. There is no need to create templates at this stage. Start sampling and capture enough data to be able to form templates, and then stop sampling. Open the Edit WaveMark dialog and use the clustering method of your choice to classify the captured data.

When you are satisfied with the clusters, use the **Apply** command to classify the original data. This will also clear all templates for the channel from the Edit WaveMark dialog and build new templates based on your classification. Now close the Edit WaveMark dialog. This will save the templates in the sampling configuration as off-line templates.

Now sample again, and when the on-line template setup dialog opens, use the **File** menu **Load and Save** command (short-cut **Ctrl+S**) and select the off-line template for your channel. This will load the templates that you created off-line.

On-line clustering We do allow you to use the Edit WaveMark dialog on-line to monitor spikes and you can also use the clustering commands on-line. However, the clustering functions will run more slowly than when working off-line; having more than one processor in your computer will help here. If you have large data files and huge numbers of events or a high data throughput to disk, this may not be a viable procedure as the system may not be responsive enough to be useful.

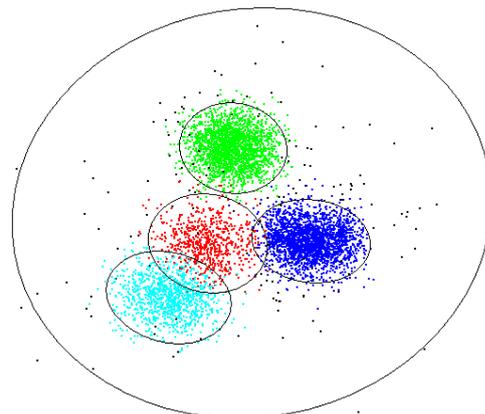
If you are using timed or triggered sampling, data sections that are not marked for writing to disk are volatile; data that the clustering code expects to find may have vanished by the time you try to apply changes.

When you apply the results of clustering on-line, the newly created templates are copied to the 1401 as if you had used the **Online update** button.

Clustering spikes

Introduction

Clustering is the generic term for methods that group similar objects into classes. In our case, the objects are spikes that consist of 1, 2 or 4 waveform traces of typically 20 to 40 data points each. The number of objects to cluster will often be large; clustering tens of thousands of spikes is not at all uncommon, so the algorithms used must be fast enough to manage large data sets in a reasonable time. Spikes also form a temporal sequence and we have to take into account that the clustering properties may change with time. The picture shows an example of four clusters and a background group.



If you have n data values that define each object, you can cluster in n dimensions. However, if n is greater than 3 it is very difficult to visualise the clusters. For spikes, n is typically 30 or more, so we need ways to extract 2 or 3 independent measurements from our data that maximise the differences we care about between the classes and minimise the differences due to noise in the data. Spike2 provides four methods: Principal component analysis, feature measurements, template correlation and template errors.

Extracting clustering values

Principal Component Analysis (PCA) is a mathematical procedure that automatically extracts the features from your data that contribute the most to the differences between the waveforms that make up your spikes. This is usually the method to try first as it is less susceptible to noise than feature measurements.

Feature measurements are user-defined values extracted from each spike, such as amplitudes, latencies, areas and slopes. Each trace is parameterized based on positions of peaks and zero-crossings. You then choose two or three measurements that emphasise the differences between the spike classes to cluster the data.

You can also cluster based on how well each spike matches two or three templates using either waveform correlations (amplitude independent) or on the sum of squares of errors.

All methods generate a table of values with one row per spike. Each table row holds the spike time, a class code and the x , y and z values derived from the spikes. All clustering operations operate on the data in this table, not on the original spikes in the data file; you can choose to apply classification changes to the original data at any time. We call the items in the cluster window *events* to distinguish them from the original spike data.

The clustering dialog contains tools that make it easy to visualise the clusters in two and three dimensions and tools that allow you to classify the clusters either manually or with various degrees of automation. How well automatic clustering works will depend on how well separated the clusters are.

Using cluster analyses

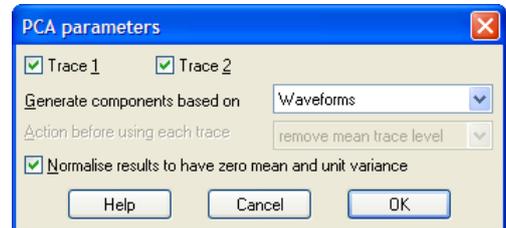
The starting point for cluster analysis is the Edit WaveMark dialog. The clustering commands are in the dialog's **Analysis** menu. The analyses operate on data in the time range set in the Edit WaveMark dialog and honour marker filters set for the data channel. Changing the Edit WaveMark channel closes any associated clustering windows.

The clustering dialog is linked to the Edit WaveMark dialog. The current event in this dialog flashes in the cluster window. If you click on an event in the cluster window, the Edit WaveMark dialog will jump to the spike that generated the event; when you apply the results of the cluster analysis, the Edit WaveMark dialog will re-evaluate its templates to match the new classifications.

Principal Component Analysis

To use this analysis, open the Edit WaveMark dialog (make sure you are not in collision analysis mode) and drag the black triangles at the top of the data display area to select the region of each spike to process. Ideally you should exclude baseline regions from the spikes to reduce noise and improve the cluster separation. The time range in the data file to process is set by the **Set Time Range** command in the dialog **Control** menu.

The **Principal Components** command in the dialog **Analysis** menu opens the PCA parameters dialog. If your spikes have a single trace, everything except the **Normalise measurements** check box is disabled. With more than 1 trace, all dialog items can be used. The items are:



Trace n When you have more than one trace you can choose which to use for analysis. If you clear all the check boxes, **Trace 1** is selected automatically. Some of the other options in the dialog require at least 2 selected traces; reducing the number of selected traces may change other dialog fields.

Generate components based on This field allows you to choose the data that the principal components are generated from when you have more than one data trace. You can choose from:

Waveforms The data source is the waveform region of each spike set in the Edit WaveMark dialog.

The data source is the peak amplitudes of the selected traces.

Amplitude at time 0 The data source is the mean peak amplitudes (ignoring sign) of the selected traces and the ratios of the peak amplitudes to the mean peak amplitude.

Mean amplitude and ratios

The peak amplitude is found by searching each trace for the peak (or trough) that is nearest to the trigger point. The trigger point is at 0.0 milliseconds in the Edit WaveMark data display window. Traces are pre-processed as set by the **Action before using each trace** field before locating the peak.

If there is not enough data to generate three output values, a z value is not generated and you will not be able to rotate the data around the x or y axis. This will happen if you choose a mode other than **Waveforms** and have 2 selected traces. It can also happen in **Waveforms** mode if the selected spike region contains less than three data points.

Action before using each trace This field allows you to decide what pre-processing should be applied to each trace before it is used. In **Waveforms** mode the mean level of the selected region of each trace is always subtracted. In other modes you can choose from:

No trace pre-processing This trace is used exactly as read from the data file.

Remove mean trace level The mean level of the trace is subtracted before using the data. You can use this to remove baseline drift.

Subtract best-fit line A straight line is fitted to the data of each trace and subtracted from the trace before the data is used. If your spikes suffer from sloping baselines you can use this method to remove them. This also removes DC offsets from your data.

Normalise measurements If you check this box, the x, y and z components generated by the principal component analysis are shifted and scaled so that they all have a mean of 0 and a variance of 1. This can make the data easier to manipulate in the clustering dialog. If you do not check this box, the displayed axes in the cluster plot originate at (0, 0, 0); events close to this point are small spikes, and are more likely to be due to noise.

Performing the analysis Click the OK button to analyse the data. This can take a long time if you have a many tens of thousands of spikes and many data points per spike. If the analysis takes longer than one second, a progress dialog appears with a **Cancel** button that allows you to abandon the analysis and return to the Edit WaveMark dialog. If a marker filter is set for the channel, only spikes that match the filter are processed.

What is Principal Component Analysis? Principal component analysis reduces a large set of non-independent data into a set of independent data that is ordered in terms of the significance of each component to the whole dataset. For most sets of spikes, only the first few components are useful, the rest correspond to noise. For clustering, the x, y and z values are the proportions of three user-selected principal components in each spike.

There is no need to be familiar with the mathematics behind principal components to make effective use of it. However, some knowledge of how it works may help you to understand its limitations and let you decide it is appropriate for your data. Principal Component Analysis is based on Singular Value Decomposition (SVD), which can be summed up by the matrix equation:

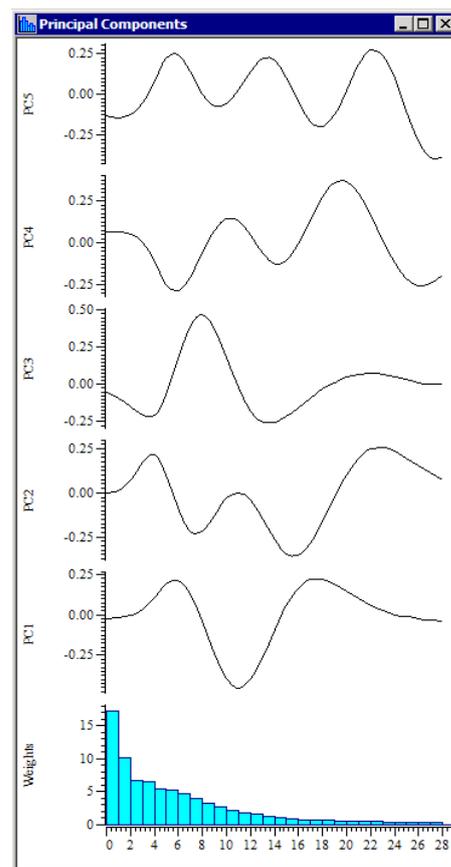
$$\mathbf{X} = \mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V}^T$$

Where \mathbf{X} and \mathbf{U} are matrices with m rows by n columns, \mathbf{W} is an n by n diagonal matrix and \mathbf{V}^T is the transpose of an n by n square orthogonal matrix. When applied to spike waveforms, m is the number of spikes, n is the number of data points in each spike. The rows of the matrix \mathbf{X} are the spike waveforms with their mean value removed. The matrix \mathbf{V} holds the principal components, the diagonal matrix \mathbf{W} holds the variance of the original data that each component accounts for and each row of \mathbf{U} holds the contribution of each component to the spike in the corresponding row in \mathbf{X} .

The principal components are ordered so that the first principal component contributes the most variance in the original data, the second represents the most remaining variance after removing the first component, and so on. If there are n points in the original spike waveform, you will get up to n principal components.

The time taken to compute \mathbf{U} , \mathbf{W} and \mathbf{V} from \mathbf{X} is proportional to mn^2 ; this means that doubling the number of spikes doubles the time, doubling the number of data points per spike quadruples the time. This is why you should not include baseline data for each spike; not only does it add noise to the analysis, it also takes longer.

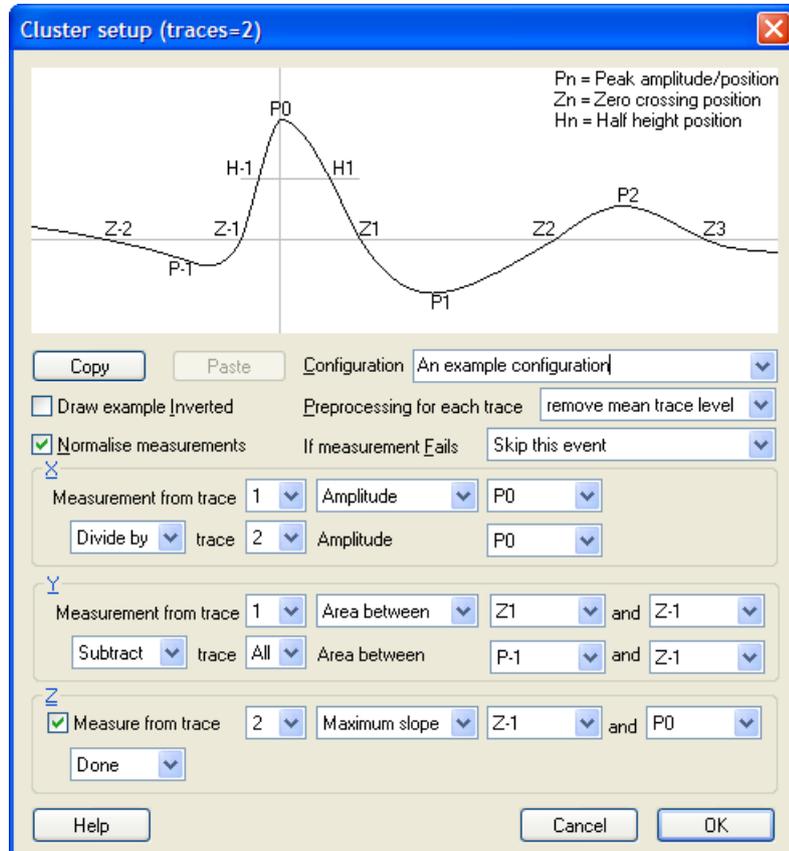
If you want to see what the components look like, the File menu Principal Components option of the clustering dialog creates a result window holding the principal components and how much each of them contributes to the original data set. In the picture, you can see the first five principal components of a set of spikes using 29 points. The Weights show the relative amplitude contribution of each of the principal components to the data. The square of the Weights is the contribution to the variance.



Cluster on measurements

When you select Cluster on Measurements from the Edit WaveMark dialog, a new dialog opens in which you can choose either 2 or 3 measurements to take from the spike data for clustering. The entire spike trace is used for measurements. The time range in the data file to process is set by the Set Time Range command in the Edit WaveMark dialog Control menu.

The dialog has 4 main regions: a schematic of a spike trace to use as a reference when selecting measurements, general controls, definitions of the x, y and z measurements, and the buttons to close the dialog. The dialog controls are:



Configuration There are 10 measurement configurations, selected with the Configuration combo box. Configurations for 1, 2 and 4 traces are stored separately. Initially, the configurations are labelled as Config 00 through Config 09, but you can edit these names to describe the type of measurement. You can copy a configuration to a different position in the table. Click the Copy button on the configuration you want to copy, select a different configuration and click Paste to overwrite it.

Draw example inverted The example waveform at the top of the diagram can be drawn either way up. This control lets you choose the direction that best suits your data.

Preprocessing for each trace If your data has a large DC offset or has a sloping baseline, measurements may not generate the expected results. You can choose to process each trace in your data before the measurements are made. You can choose from:

No trace pre-processing This trace is used exactly as read from the data file.

Remove mean trace level The mean level of the trace is subtracted before using it.

Subtract best-fit line A straight line is fitted to each trace and subtracted from it before the data is used.

If measurement fails It is not always possible to make a measurement. For example, a peak might not exist in the data. You can choose between:

Use standard work-around This sets un-measurable amplitudes to 0 and un-measurable times to a time just past the end of the spike when searching forwards or to a time just before the start of the spike when searching backwards.

Skip this event The spike is not included in the measurements or clustering. If you use the **Apply** command in the cluster dialog, all skipped spikes are given code 00.

Normalise measurements If you check this box, the measurements for x, y and z are shifted and scaled so that they all have a mean of 0.0 and a variance of 1. This can make the data much easier to work with in the cluster dialog, especially if you rotate the data. Of course, the measurements are then in arbitrary units, so you may not wish to do this if you want to export the values as text from the cluster dialog.

X, Y and Z measurements These three areas are identical except that you can enable and disable Z measurements with the **Measure from trace** checkbox.

We model each trace by searching outwards from time 0 for the peak P0. We expect this peak to exist because the spike data was aligned on a positive or negative peak when it was captured. Next we search outwards in both directions from P0 to find the zero crossings Z-2, Z-1, Z1, Z2 and Z3. Then we find the peaks P-1, P1 and P2. Finally we locate the half peak height positions H-1 and H1. To make the searches as accurate as possible, we fit a cubic spline through the trace data points and search for peaks and level crossings using the fitted data. In addition to these measurements, there are also:

PkMax The position and amplitude of the highest (most positive) peak.
TrMin The position and amplitude of the minimum (most negative) trough.
Max The position and amplitude of the most positive point in the trace.
Min The position and amplitude of the most negative point in the trace.

The top line of each area sets a basic measurement. The first field on the second line can be set to one of **Done** to use this value or **Subtract** or **Divide** by to make a second measurement and subtract this from the first or divide the first measurement by the second. You can choose from the following measurements:

Amplitude	You can select any of the peaks (P-1 to P2) as the measurement. The units of measurement are the y axis units of the channel.
Time at	The measurement is the time of any feature (Peak or level crossing). The units are milliseconds relative to the peak trigger point (the 0 time in the Edit WaveMark dialog).
Area between	The measurement is the area between any two features, from the curve to the y axis zero. The units of the area are y axis units times milliseconds. All areas are treated as positive.
Maximum slope	The measurement is the value of the steepest slope of the fitted cubic spline between any two features in y axis units per millisecond.
Best fit slope	The measurement is the slope of the least-squares best-fit line between any two selected features in y axis units per millisecond.

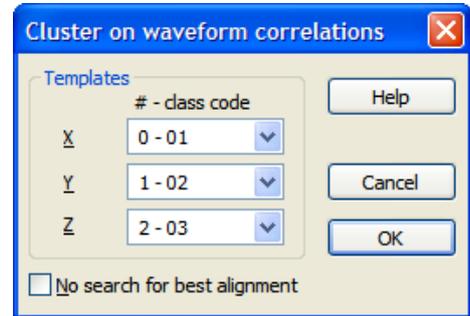
If you are working with multiple trace data you can choose which trace to use as a source of your measurement. If any part of the measurement calculation uses a failed feature measurement, the measurement is marked as failed. You can also choose **All traces**, in which case the measurement is the average measurement value on all of the traces.

Running the analysis Click on the **OK** button to save any changed configuration settings and generate the measurements. The analysis process takes a time proportional to the number of spikes. Once the analysis is done the Cluster dialog will appear. The **Cancel** and **OK** buttons

change to Close and Apply if you open this dialog with the Reanalyse command from the clustering window.

Cluster on template correlations

To use this option, click on the Cluster on Correlations command in the Analysis menu of the Edit WaveMark dialog. There must be at least 2 templates defined in the dialog for this command to be enabled. The command opens a dialog in which you can choose two or three templates to correlate each spike with to generate the X, Y and optionally Z values for clustering. You can choose NONE as an option for the Z value.



No search for best alignment

If you check this box, the section of the spike shape correlated against each template will be the section marked for template formation in the main display in the Edit WaveMark dialog. If you leave the box unchecked, Spike2 will search around the section marked for template formation for the best correlation. This search extends for up to two sample points in either direction and uses cubic spline interpolation of the spike waveform to estimate values between sample points.

You will normally leave the box unchecked unless you have chosen a template region that does not include the spike alignment point (usually the peak or trough), and the time delay of similar shapes is the distinguishing criterion. If you check the box, the analysis is much quicker, but the result is usually significantly worse.

How the correlation is calculated

The correlation between a segment of the spike shape starting at position s and a template of n points is calculated as follows:

- 1) Make a spike segment of n points (use cubic spline interpolation if s is not integral).
- 2) Normalise the spike segment by shifting and scaling so that the mean value is zero and the sum of squares is unity.
- 3) Normalise the template to have zero mean and unity sum of squares.
- 4) The correlation is the sum of the point by point product of the two waveforms.

The result of this is a value between 1.0 (if the spike and template have the same shape) and -1 (if they have the same shape but one is inverted with respect to the other). If the two shapes are significantly different, the result will be nearer to 0 than to 1 or -1.

When to use correlations

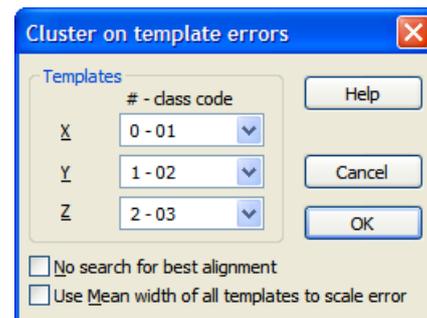
The correlation has the property that it is *independent of the amplitude of the spike*; the correlation is purely a function of shape. As spike amplitude is often the single most important feature when sorting, you should consider carefully whether this is worth losing. On the other hand, if you have a spike of more or less constant shape, but variable amplitude, this may be exactly what you want.

We anticipate that this method will be useful in sorting similar spikes with varying amplitudes into two or three categories, possibly after using other sorting methods to reduce the problem down to the difficult cases.

The results produced in three dimensions tend to be a thin layer of events with patches at a radius of 1 from the origin for each of the templates, which may not be ideal for the automatic clustering. You may find that you have to use manual event selection.

Cluster on template errors

To use this option, click on the **Cluster on Errors** command in the **Analysis** menu of the **Edit WaveMark** dialog. There must be at least 2 templates defined in the dialog for this command to be enabled. The command opens a dialog in which you can choose two or three templates to compare each spike with to generate the X, Y and optionally Z values for clustering. You can choose **None** as an option for the Z value.



No search for best alignment

If you check this box, the section of the spike shape compared with each template will be the section marked for template formation in the main display in the **Edit WaveMark** dialog. If you leave the box unchecked, **Spike2** will search around the section marked for template formation for the smallest error. This search extends for up to two sample points in either direction and uses cubic spline interpolation of the spike waveform to estimate values between sample points.

You will normally leave the box unchecked unless you have chosen a template region that does not include the spike alignment point (usually the peak or trough), and the time delay of similar shapes is the distinguishing criterion. If you check the box, the analysis is much quicker, but the result is usually significantly worse.

Use Mean width of all templates to scale error

With this unchecked, the error at each comparison point is scaled by the template width at that point. If you check the box, the error is scaled by the mean width of all templates (with multiple traces, the error for a trace is scaled by the mean template error for that trace in all templates). If you check this box, the error can be thought of as the least squares error. If you don't check it, the error is more like a chi-squared error based on the template width. Use whichever gives the better separation.

How the error is calculated

The error between a segment of the spike shape starting at position s and a template of n points is calculated as follows:

- 1) Make a spike segment of n points (use cubic spline interpolation if s is not integral).
- 2) Adjust the spike segment to have a mean value of zero.
- 2) Adjust the template to have zero mean.
- 3) Form the sum of the squares of the point-by-point difference of the two waveforms divided by either the point-by-point width or by the mean width.
- 4) Divide the result by the number of points to form the average error per point.

The result of this is a value from 0 upwards. You would expect values for a matching template to be less than 1.

When to use errors

We anticipate that this method will be useful in sorting similar spikes into two or three categories, possibly after using other sorting methods to reduce the problem down to the difficult cases. Spikes that resemble each of the templates should lie in the range 0 to 1 on the axis that corresponds to the template.

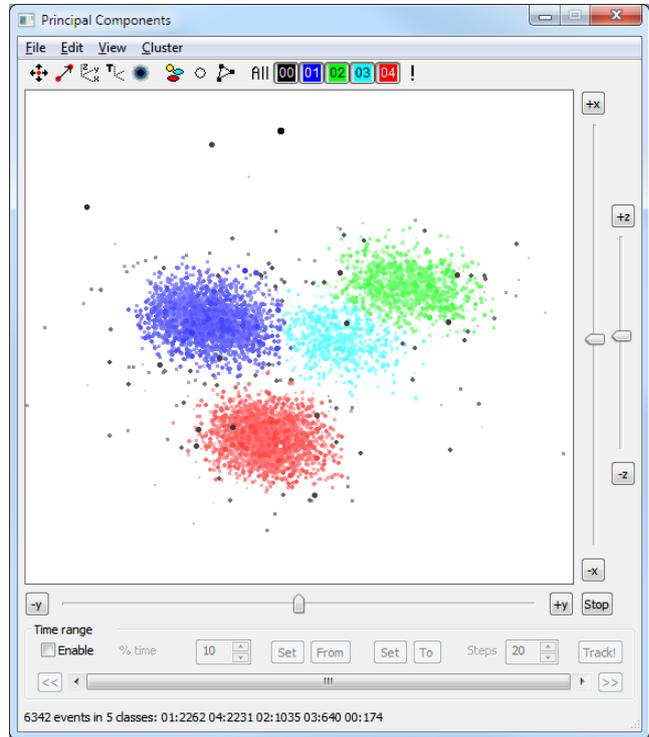
If you check the *Mean width* box, this method is similar to the method used for template matching in the **Edit WaveMark** dialog and you can use it to get a visual indication of the reliability of the separation achieved by the template method. If this produces non-overlapped clusters, it is likely that the template method is reliable.

The Clustering dialog

You reach this dialog by using one of the clustering commands in the Analysis menu of the Edit WaveMark dialog. The dialog behaves almost identically for all analysis methods; the dialog title indicates the source of the data. The dialog contains a menu, a toolbar that selects what to display, the cluster window, sliders for rotation around the x, y and z axes, the Time range area and a status line. You can hide the rotation slider controls and the Time range area.

You can resize the dialog by clicking on any edge or corner. The minimum size depends on the displayed items; if you hide the Time range area and the rotation controls you can make the dialog very small indeed!

The analysis you selected built a table of events. Each event contributes one dot to the display. For each event, the table holds the associated spike time, the x, y and z values from the analysis and the class code. The initial class codes in the table are copied from the spikes. Each spike has 4 marker codes; we use the marker code selected in the Edit WaveMark dialog. Class assignments made in this dialog only change the table. The File menu Apply changes command copies changes to the data file. There are two methods you can use to display the events:



Density plot Select this display mode with the View menu Density Plot command. The display shows the density of events in each pixel. The View menu Density Colour Map command lets you change the colour scale used to indicate the event density. The View menu Density Settings command gives you control over the creation of the density map.

Class code colours In this display mode, the colour of each dot is set by the class code (matching the WaveMark colours used in the time window) and the background colour for the cluster region is set in the main Spike2 colour palette. If the event corresponding with the spike displayed in the Edit WaveMark dialog is in the visible area, it blinks on and off.

Class ellipsoids We calculate statistics for each class of events and based on the assumption that the event x, y and z values in each cluster are normally distributed about the cluster centre, we calculate ellipsoids of constant probability around the cluster centres. You can set the size of the ellipsoids in terms of the Mahalanobis distance (this is multi-dimensional version of standard deviation) with the View menu Ellipse radius command. We display the projection of the ellipsoid onto the cluster window.

User shapes We also maintain a two-dimensional user ellipse and a user-defined shape that are used to surround events and give them a user-defined class. The user ellipse is independent of the three dimensional class ellipsoids.

Toolbar controls There is a Toolbar across the top of the display. The buttons are arranged in three groups:



The group on the right controls the display of events in the cluster window. There is one button for each class code that is present in data. Each button displays an event colour and class code. The colour scheme is the same as for WaveMark data in a time view. Events are displayed for buttons that are down, events are hidden when buttons are up. Only visible events are modified by the commands in this dialog. The All button is a short-cut to display all the event codes. Right-click All to hide all codes. The ! button inverts the displayed class codes; all visible event classes are hidden, all hidden event classes are made visible.



The group on the left provide short-cuts for various menu commands. You will find the details of these commands in the View menu description, below. From left to right:

- Autoscale Scale the image to keep everything visible.
- Link close events Link events closer that a minimum time period with an arrow. Right click this button to open the Minimum Interval dialog.
- Show axes Display x, y and z axes in the display area.
- Use Time as Z Use event times as the z axis for display purposes
- Density plot Display a density plot. Right-click to open the settings dialog.



The central group controls the display of shapes that either describe the extent of a class of events or that are used to mark events as belonging to a class From left to right:

- Class ellipses Display ellipses for each visible class of events. Right-click the button to open a dialog that controls the size of the ellipse.
- User ellipse Display a user-defined ellipse. This is used to select events for coding and to restrict the events used in the interval histogram.
- User shape Display a user-defined shape formed from straight lines. This is used in the same way as the User ellipse.

Identifying events

The mouse pointer changes to a cross hair when it is over the cluster window. To find out which event produces a particular dot in the window, centre the cross hair on the dot and click the mouse button and release. Spike2 will locate the nearest event to the centre of the cross hair and display it in the Edit WaveMark dialog. Cursor 0 will also move to the spike position in the associated time view. This does not work if you are on-line and the *always take spikes from the end of the file* button is down in the Edit WaveMark dialog.

Scaling and shifting the cluster window



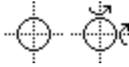
You can zoom in to an area in the cluster window by clicking and dragging with the mouse. While the mouse is down, the mouse pointer changes to a magnifying glass with a plus sign in the centre and a dotted rectangle shows the selected area. When you release the mouse, this area expands to fill the window. If you hold down the `Ctrl` key before you click and drag, the magnifying glass has a minus sign in the centre and when you release the mouse, the original area of the window is scaled to fit in the dotted rectangle.



You can drag the cluster window by holding down both the `Shift` and `Ctrl` keys before you click and drag the display. The mouse pointer is an open hand before you click and a closed hand when you drag.

Rotating the cluster window

There are three sliders with associated buttons that rotate the cluster window around the x, y and z axes. Initially, x runs from left to right, y runs from bottom to top and z is out of the screen towards you. The small buttons at the ends of the sliders (labelled +x, -x, +y, -y, +z and -z) set a continuous rotation. Each time you click one it changes the rotation rate. The Stop button cancels continuous rotations. If the sliders are not visible you can use the `Ctrl+S` short-cut key to show them.



You can also rotate the cluster window using the mouse. If you hold down the `Shift` key and move the mouse over the cluster window, the mouse pointer changes to a cross-hair with a circle. Now imagine that the events are in a large glass sphere and that you are looking at it through a square window. If you click and drag, with the `Shift` key down, this "grabs" the surface of the sphere and moves it in the direction that you are dragging.

When you rotate the cluster window, any displayed class ellipses also rotate. If the user ellipse is present it does not change.

In some cases, the data resulting from the analysis will not have any `z` information. In this case, rotation with the `x` and `y` sliders and by clicking and dragging is disabled. You are still allowed to use the `z` slider to rotate the data around the `z` axis.

Jitter the cluster window

Instead of rotating the view, you can apply a small circular jitter to the display with the `Jitter` command in the `View` menu. This allows you to rotate the display to best view the clusters, then apply a small jitter to help you visualise the clusters in three dimensions.

Working with ellipses



You can select, drag and modify the shape of displayed ellipses. Click on or within an ellipse to select it; four resizing handles appear to show that it is selected. To select multiple ellipses, hold down the `Ctrl` key before you click. If you click in a region that is enclosed by more than one ellipse, the ellipse with the smallest area is selected.

To move a selected ellipse, click inside the ellipse and drag. To resize an ellipse, click on one of the handles and drag. The opposite handle is fixed while you drag. If you hold down `Ctrl` and drag, the ellipse resizes around its centre. If the display window scaling is different in the `x` and `y` directions you will find that the ellipses distort if you rotate them by dragging a handle. You can force the scales to be the same in both directions with the `View` menu `Equal scales` command.

When you select one or more ellipses, the status bar displays a list of the classes and the number of events in each class that lie within the selected areas. The class list is sorted with the largest class first.

You can change the class codes of the events inside or outside selected ellipses with the `Cluster` menu `Set Codes` command. As a short-cut, if you press the keys `0` through `9` with one or more ellipses selected, all events within the selected ellipses change class to match the key: `0` for class `00` up to `9` for class `09`. You can also set the codes of all visible events with the `Ctrl+Shift+0` through `Ctrl+Shift+9` key combinations.

Although you can change the positions and sizes of the class ellipses, such changes have no effect on the classes unless you use one of the methods discussed above to change the class codes. To restore the positions of the class ellipses, click the toolbar `SD` button twice: once to hide ellipses, and a second time to restore them to their original positions.

Working with the user-defined shape



If your clusters are not ellipsoidal, you may find it easier to select items with the user-defined shape. To start a shape, hold down `Alt` and click in the cluster window. Then move the mouse and click to add additional corners. To close the shape, click at or near the start point, or double click to add a point and then close the shape.

The user-defined shape behaves in much the same way as the user ellipse. You can select it and move it around the cluster window. You can select individual handles and drag them to adjust the positions of the corners. With the shape selected you can set classes with the `Cluster` menu `Set Codes` command, exactly as for the user ellipse.

Menu commands The clustering dialog has its own menu. All references to menu items in this section of the documentation are to the dialog menu, not the main Spike2 menu. Most menu items have short-cut keys that allow immediate use; you can also right-click in the dialog to open a context menu that duplicates many of the menu commands.

- File menu** Apply changes, restore the original class codes from the data file, reanalyse the data, create interval histograms of the data and display the principal component waveforms.
- Edit menu** Undo class changes, copy an image, copy data as text and edit interval histogram settings.
- View menu** Everything to do with the appearance of the display and control over rotation, automatic scaling and ellipse and dot sizes.
- Cluster menu** Commands for automated data clustering and setting class codes.

File menu This contains the following commands:

Apply changes This copies the current state of the classes in the clustering window into the data file. Any events that were skipped during measurement analysis are treated as having a class code of 00. For this to work as expected you must not have changed the marker filter for the data channel since extracting the clustering values. If you do not use this command, all class changes you make in this dialog are lost when the dialog closes.

In addition to changing the events in the data file, the associated spike shape dialog will recalculate its templates based on the class assignments of the events in the marker filter and current time range set in the dialog. If you are working on-line, the newly created templates are copied to the 1401 for immediate use.

Restore codes This command sets all the class codes to match the data file. You can use this to restore values after an unsuccessful clustering attempt. Another use is to pass class information between the measurement and PCA clustering dialogs if you have both open at the same time; use **Apply changes** in one and **Restore codes** in the other.

Reanalyse This command opens the Principal Components, Cluster on measurements, Cluster on Correlations or Cluster on Errors setup dialog so that you can repeat the analysis.

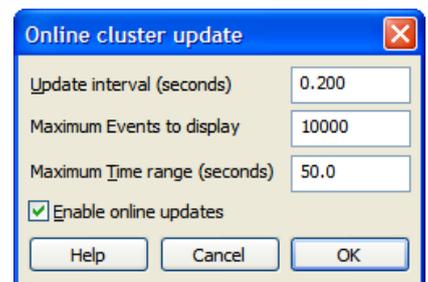
Online Update This File menu command is enabled if you are clustering a channel that is being sampled. There are four fields you can set:

Update interval This sets the minimum gap between updates of the cluster window in seconds. A value of 0 updates the window as often as possible.

Maximum events to display This limits the number of events to display in the window. At high event rates, if you set a large number and a short update interval, Spike2 may become slow to respond as it is spending a lot of time recalculating the display.

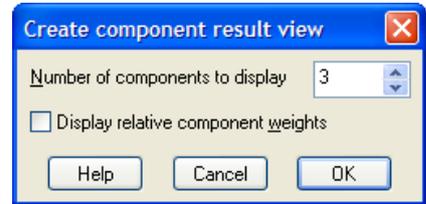
Maximum Time range This field limits how far back from the current time to search for events. The combination of this and the *Maximum events to display* field limits the number of displayed events.

Enable online updates You can enable and disable online updates here. If you disable them, the Reanalyse command will process the time range set for the Edit WaveMark dialog, or the entire file if the Edit WaveMark dialog is in At End mode.

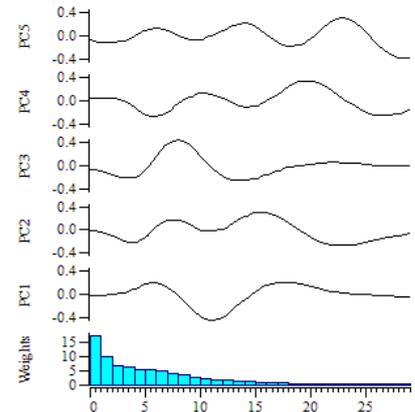


Interval histogram This command creates an interval histogram of the events in the clustering window. If an interval histogram of the events already exists, the previous values in the result view are replaced by a new analysis. Each class code present in the clustering window generates a separate channel in the histogram. The histogram bins are coloured to match the class colours except for code 00, which is filled with white rather than black. You can set the histogram width and bin width with the Edit menu INTH Settings command. If you display a User ellipse or User shape and select it, only events that lie inside the selection appear in the INTH.

Principal components This is enabled for Principal Component Analysis. It creates (or updates if you have already created it) a result view holding the principal component waveforms and optionally the relative weights of each component. You can choose the number of components to display. The display order is in terms of their significance, based on their contribution to the variance of the original data.



If you choose to display the weights, the x axis of the result view is the component number for the weights which corresponds with data points for the principal component waveforms. If you do not display the weights, the x axis is set to milliseconds with zero time representing the trigger point when the data was captured. If the spike data has multiple traces, each component shows all the traces, in order.

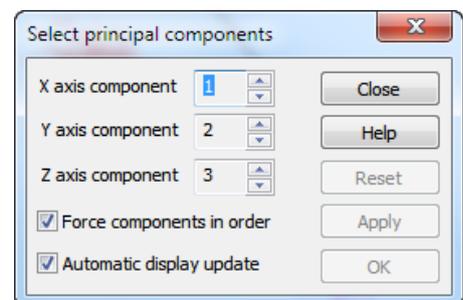


The above assumes a waveform-based analysis. Analyses of amplitudes or amplitude ratios will not make any sense when drawn as waveforms.

Select Principal Components This dialog sets the principal components to display. Reset selects components 1 to 3; these are usually the best ones to choose.

Force components in order Check this box to make sure that the X component has the lowest number and the Z component has the highest.

Automatic display update Check this to update the display on every change; this can be slow if you have a very large number of events. The Apply button can be used if you disable automatic updates.



Close This closes the dialog. It does not save any changes you have made to the spike classes.

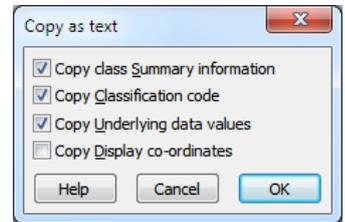
Edit menu This contains the following commands:

Undo This undoes the last change made to the class codes. At the time of writing only one change is saved. In the future we may allow more than one undo.

Delete online spikes This command and the shortcut key Del are enabled when you are working online. It deletes all the spikes from the cluster.

Copy This command copies the cluster window to the clipboard as a bitmap. The bitmap image is always created with z buffering enabled, so events that are further away from the viewer are hidden by events that are closer. This may differ from the screen image where the z buffer is only used if selected because it slows down drawing.

Copy As Text This menu command opens a dialog from which you can copy clustering data to the clipboard as text in a spreadsheet-compatible format. The output is in columns separated by tabs, with text enclosed in double quotes. There is one line of output for every visible event in the cluster window. The output has the form:



```
6342 events in 3 classes: 01:3300 02:3002 00:40
"Time" "Code" "X" "Y" "Z"
0.00016 "02" -0.0107435 0.00288545 -0.0126408
0.17096 "02" -0.00890594 0.00563056 -0.0174058
```

The time of each event is always output in the first column. You can also choose to copy the following, which are described in the order that they are added to the output line:

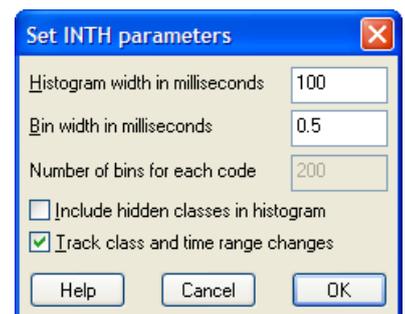
- Summary information* This is the same as the line of text displayed at the bottom of the cluster window.
- Classification code* Check this box to output the class code for each event, with the column title "Code".
- Underlying data* The x, y and z values from the analysis or measurement. The columns titles are "X", "Y" and "Z". If there are no z values, the z column is omitted.
- Display co-ordinates* These are the underlying data values rotated into display co-ordinates. The x and y values set the event positions in the cluster window. The z value is used for the z buffer. The titles are "View X", "View Y" and "View Z". If there are no z values in the original data, the "View Z" column is omitted.

Copy Cluster Values This menu command copies information for all visible cluster classes to the clipboard as text. Tab characters separate the columns of text and the titles are surrounded by double quotes. The output columns contain: the cluster code, the number of items in the cluster, then for each axis the cluster centre and the sigma value (square root of the variance about the centre) and the mean cross products about the centre (these would be 0 if the clusters were aligned with the axes). The Z values are omitted if there is no Z axis.

```
"Cluster" "N" "X" "SigmaX" "Y" "SigmaY" "Z" "SigmaZ" "XY" "XZ" "YZ"
00 17 -1.4703 0.9867 0.2347 0.9501 0.0469 1.0200 -0.434 0.1411 0.49357
01 2348 -1.0189 0.4626 0.4006 0.2660 0.2213 0.8000 0.000917 0.00818 0.01794
02 1036 1.4826 0.4147 1.3981 0.2884 -0.3296 1.1811 -0.001983 -0.0379 0.10251
03 645 0.7424 0.3856 0.4463 0.2788 -0.2948 1.0093 -0.005489 -0.02464 0.04993
04 2218 0.1815 0.3615 -1.2086 0.2790 0.0050 1.0308 -0.001484 -0.00329 0.08619
```

In this example, the mean cross products are small compared to the variance (apart from cluster 00), meaning that the data values for the three axes are independent. This data came from a principal component analysis, so we would expect this to be the case.

INTH Settings This menu command opens a dialog that sets the parameters for creating an interval histogram based on the events in the cluster window. The histogram has multiple channels, one for each event class. When you close the dialog with the OK button, the interval histogram is created or updated if it already exists. The File menu Interval histogram command also uses the values set in this dialog.



You can set the desired histogram width and the width of each bin, both in milliseconds. The number of bins in the histogram is also displayed, but this field cannot be edited.

If you set the **Include hidden classes in histogram** checkbox, all event classes are processed and displayed. If the checkbox is clear, only visible classes are included and the INTH changes when you show or hide classes in the cluster window. If you check the **Track class and time range changes** checkbox, any operation that changes the visible events or the event classes causes the interval histogram to recalculate.

View menu This menu (plus suitable entries depending on your current selection) can also be activated as a context menu by right-clicking the mouse in the cluster window. The menu contains the following commands:

Autoscale When you rotate the cluster window you may find that the data tends move out of the viewing area. Turn this option on to rescale the view to fit all visible events, axes and class ellipses in the window. The view rescales automatically if you change the displayed classes, the displayed time range or the window rotation.

If your interest is purely in the clusters and not at all in the actual value used for clustering, you can often remove the need for scaling by checking the **normalise** box in the setup dialog for the analysis you used to create your cluster data. This forces the x, y and z values to have a zero mean and a unity variance.

Equal scales If you enable this option, the cluster window is scaled so that the width and height of the window have the same numeric size. This makes ellipses easier to manage, as they will preserve their shape when rotated. The more different the scaling of the x and y display axes, the more the ellipses will distort when rotated.

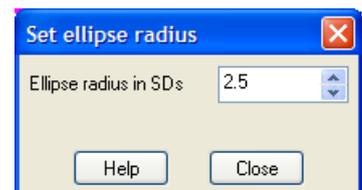
Z Buffer With this option on, events nearer to the viewer draw on top of events that are further away. This increases the illusion of three dimensions, but may draw more slowly. When you copy the image to the clipboard, this option is always turned on. The previous state is restored once the bitmap has been copied.

Dot size The events in the cluster window can be drawn at five sizes, numbered 0 to 4 or **Varies** for a dot size that depends on the Z position of the event in the display.

Colour Fade with Z This is most effective with the Z buffer enabled. When enabled, the colour of each event fades towards the background colour as the events get further away from the viewer.

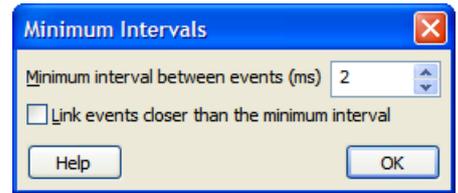
Circular Dots You can choose to draw circular dots to represent the events. This is only effective at dot sizes 3 and 4, and looks best at size 4.

Ellipse radius This command opens a dialog in which you can set the size of the class ellipses that are drawn around each class cluster. The size is set in terms of the Mahalanobis distance, which is a multi-dimensional analogue of the standard deviation. If the data were normally distributed around the mean position, you would expect 68% of the events to be within the Mahalanobis distance, 95% of the events to be within twice this distance and 99% to lie within three times this distance.



You can set the distance to be any value in the range 1.0 to 4.0 times the distance. The default value is 2.5 times, which seems to be about right. If the class ellipses are visible, any change made in the dialog has immediate effect.

Minimum Interval One way to check that codes set by clustering are reasonable is to look at the time intervals between events in the same class and verify that events are not too close together. The Minimum Interval command opens a dialog with the following fields:



Minimum interval between events (ms) This field sets the minimum acceptable interval, in milliseconds, between events with the same class code.

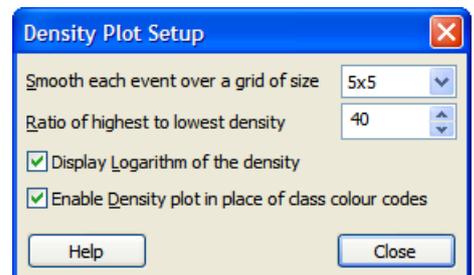
Link events closer than the minimum interval If you check this box, pairs of events that are closer than the minimum interval field are linked by a line with an arrow pointing from the earlier event to the later event. You can click on the events to move to them in the Edit WaveMark dialog. Click near an arrow or right-click within 100 pixels and select Find Short Interval to locate the event that starts the nearest short interval to the mouse pointer. There is a toolbar short-cut for this action.

Use Time as Z You can display events using Time as the Z axis. This is not useful when clustering. Event times are mapped into the range -1 for the oldest to +1 for the newest. Applying and removing this option can be slow with a large number of events. This option is not allowed if you are in online update mode, or if you have had an online update.

Rescale This command scales the display window so that all the displayed event data and any class ellipses are visible. If the Equal scales option is set, the larger of the two display axes sets the scaling for both axes. If axes are enabled, the axis origin is also displayed.

Density Plot There are two basic display modes: Density Plot and Class Colours. This command chooses between the two modes. In the Class Colours mode, each event is drawn as a dot in the colour of the event class. In Density Plot mode, each pixel of the display is drawn in a colour that shows then number of events at that point in the display.

Density Settings This command leads to the Density Plot Setup dialog, which controls how the event density is converted into a density display.



Smooth each event over a grid of size To convert the events from individual points to a smoothly varying density plot, each event is spread over a square grid of pixels that is centred on the event and the result of this is summed over all the events. You can choose from a 1x1 grid (no smoothing) up to a 7x7 grid.

Ratio of highest to lowest density The highest density is allocated one end of the density colour scale. This figure sets the lowest visible density relative to the highest. Smaller numbers will make individual outlier events merge into the background.

Display Logarithm of the density If all your clusters have similar densities, then it makes sense to map densities to your colour map linearly. However, if you have clusters of widely differing densities, it can make more sense to display the logarithm of the density.

Enable Density plot in place of class colour codes This is exactly the same as the Density Plot command and allows you to swap between a display of class colours and density from within the dialog.

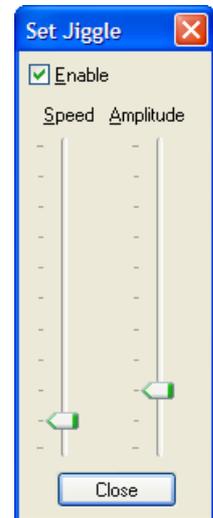
Density Colour Map This command allows you to choose the colour scale that is used to indicate density. This dialog is exactly the same as for the sonogram colour scale described for the Edit menu Preferences command.

View along axis This command (enabled with three dimensional data only) display a pop-up menu with three options: X, Y and Z. Choose one of these options to rotate the cluster view to point one of the three data axes upwards and cancel any display rotation. It does not cancel any display jiggle. The keyboard shortcuts for these commands are X, Y and Z.

Jiggle Display and Jiggle Settings

With three-dimensional data, it is sometimes easier to visualise the clusters if the display is "jiggled" around the current view angle. If you enable the display jiggle, this is equivalent to using the mouse to rotate the display in a circle around the centre of the displayed cluster area. You can enable and disable this rotation with the Display Jiggle command or with the `Ctrl+J` keyboard shortcut.

The Jiggle Settings dialog gives you additional control over the jiggle. You can set both the amplitude of the jiggle (equivalent to the radius of the circle) and the rate of rotation. The **Enable** checkbox starts and stops the display jiggle. The keyboard shortcut to open the dialog is `Ctrl+Shift+J`.



Clear rotation This command restores the cluster window to its original, non-rotated state and removes any continuous rotation set by the buttons around the rotation slider controls. It also cancels any display jiggle.

Show Sliders You can choose to show or hide the slider controls and associated buttons that rotate the clustering window around the x, y and z axes. The dialog size will grow, if required, when you turn on the sliders.

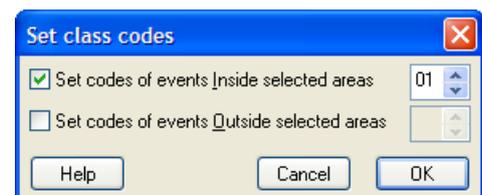
Show Time Range The Time range control area can be shown and hidden as required. The dialog will change size when you show and hide this area. This area allows you to display events that fall within a restricted time range and track how the clusters change with time.

Show Axes You can display axes in the cluster window. They are labelled x, y and z and are drawn along the principal axes of the measurements.

Cluster menu This menu contains commands related to manual and automatic assignment of class codes to the clusters of events. The menu commands are:

Set codes

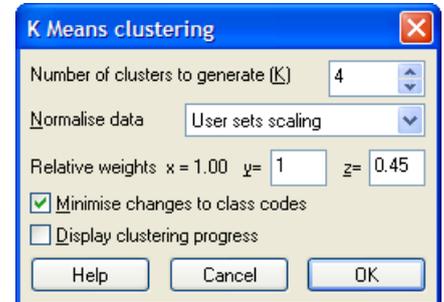
This command is enabled when there are selected shapes in the cluster window. It opens a dialog in which you can set the codes of all events that lie inside and/or outside selected shapes. When you open the dialog the inside code is set to the most common code of events inside selected shapes unless this code is 00, in which case the lowest unused class code is set. The outside events code remembers the last code you set.



The check boxes set which events to change. The edit boxes on the right set a code to apply as two hexadecimal digits (0-9, a-f or A-F). If you set a single character in the field, a 0 will be inserted in front of it. Leaving the field blank is the same as clearing the check box and no class codes will change. Click **OK** to change the displayed events. This has no effect on the original spikes in the data file. You must use the **File** menu **Apply** command to change their class codes. You can undo the effect of this command with `Ctrl+Z`.

Clear all visible codes This command is here as a quick way to set the class codes of all visible events to 00. It has a short-cut key combination `Ctrl+Shift+0`. In fact, you can set all visible events to any code from 00 to 09 using `Ctrl+Shift+0` through `Ctrl+Shift+9`. You can undo the effect of these commands with `Ctrl+Z`.

K Means from existing and K Means These commands open a dialog in which you can assign events to classes automatically. The **K Means from existing** command uses the existing cluster centres as the seeds for the K Means algorithm. The **K Means** command runs the algorithm 10 times (or until you stop it) with random seeds and chooses the result that has the best ratio of cluster separation to cluster size. You can find details of the algorithm at the end of the chapter.



The K Means algorithm is a well-known and fast clustering method that is effective when the number of clusters is already known, the clusters are spherical and each cluster holds a similar number of events. In our case, you can probably guess the number of clusters, but they may not be spherical and they certainly won't have the same number of events!

We can often make the data spherical by scaling the data. This works when the principal axes of all the class ellipsoids are aligned with the x, y and z directions. This is often true for principal component data, but is less likely for data derived from measurements. If some clusters contain few events compared to others, the algorithm may ignore smaller clusters and prefer to split clusters with many events. You can work around this by classifying the larger clusters first and hiding them before working on the smaller ones.

Number of clusters to generate This field sets the number of clusters to generate. It is preset to the number of visible, non-zero class codes in the data set. We allow you to set from 1 to 20 clusters. This field is disabled if you have selected the **K Means from existing** command.

Normalise data This field sets the strategy to use to make the clusters spherical. You can choose from:

None	No scaling: using the measured values directly. This is equivalent to the User sets scaling option with all weights set to 1.
Use existing clusters	The relative scaling of the x, y and z data is deduced from the shape of the current class clusters. As clustering changes the class clusters, you may need to run this several times to get a stable result.
Use visual scaling	We assume that you have arranged the display so that the clusters appear circular/spherical when you rotate the display. The data is clustered based on the visual appearance of the data in the window.
User sets scaling	Extra fields appear in which you can set the relative weights to give to the y and z values relative to the x values. For principal component analysis, these fields are initially set to values based on the relative weights of the components. For measurement-based clustering, the relative weights are initially set to 1.

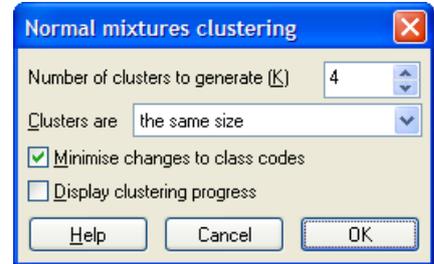
Minimise changes to class codes If you check this box, the event class codes are chosen to maximise the number of events that keep the original code. If you do not check this box, class codes are assigned so that the lowest free codes are used. This is always checked for **K Means from existing**.

Display clustering progress Check this option to follow the clustering in detail. The **K Means** command updates the display after each iteration regardless of this setting. This option slows analysis.

Running the command Click OK to run the command. This may take some time if there is a large number of events and K is large. When the analysis is complete the command calculates the ratio of the sum of squares of the distances of all points from centre of all the points to the sum of squares of the distances of all points from the centre of the cluster they belong to. This is displayed in the status line as the J3 value.

Normal Mixtures from Existing and Normal Mixtures

These commands open the Normal Mixtures dialog in which you can assign events to classes automatically. The Normal Mixtures from existing command uses the existing clusters as the seeds for the Normal Mixtures algorithm. The Normal Mixtures command runs the algorithm 10 times (or until you stop it) with random seeds and chooses the result that has the maximum likelihood of fitting the data.



The Normal Mixtures algorithm assumes that the probability density of data points around each cluster centre follows a multivariate normal distribution. This is a more general approach than the K Means algorithm, as it does not require spherical clusters and can often give results that seem more intuitively correct. However, it is a more complex algorithm than K Means and takes noticeably longer to run.

Number of clusters to generate This field sets the number of clusters to generate. It is preset to the number of visible, non-zero class codes in the data set. We allow you to set from 1 to 20 clusters. This field is disabled if you have selected the Normal Mixtures from existing command.

Clusters are You can use this field to place restrictions on the clusters. The choices range from least to most restrictive. You can choose from:

- Independent** The clusters are independent of each other and have different sizes, shapes and orientations. This is the most general setting and takes the longest to run. It is usually most successful when run with existing clusters as the starting point. Using random starting points can generate unexpected results.
- the same size** The clusters are assumed to be the same size and shape. This is often true, especially with principal component data, so this is the default value.
- the same size and axially aligned** The clusters are assumed to be the same size and shape, and the principal axes of the cluster shapes are aligned to the x, y and z axes of the data.
- the same size and spherical** The clusters are assumed to be the same size and shape and are spherical. This is the most constrained setting.

Minimise changes to class codes If you check this box, the event class codes are chosen to maximise the number of events that keep the original code. If you do not check this box, class codes are assigned so that the lowest free codes are used and ordered so that classes with more events have lower codes. This is always checked for the Normal Mixtures from existing command.

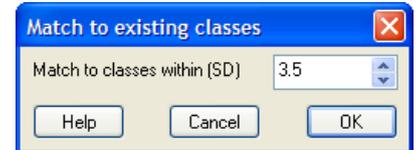
Display clustering progress Check the box to follow the clustering in detail. The Normal Mixtures command updates the display after each iteration regardless of this setting. This option slows analysis.

Running the command Click OK to run the command. This may take some time if there are many events and the number of clusters is large. During analysis, a progress bar gives you the option to stop

analysis early. When the analysis is complete the command displays the log likelihood value for the solution.

Match to classes

This command opens a dialog from which you can use the current cluster statistics to assign visible events to the nearest visible cluster and mark outliers as code 00. The typical use of this command is to clean up manual cluster assignments in situations where the K Means from Existing or Normal Mixtures from Existing commands would make too drastic a change.

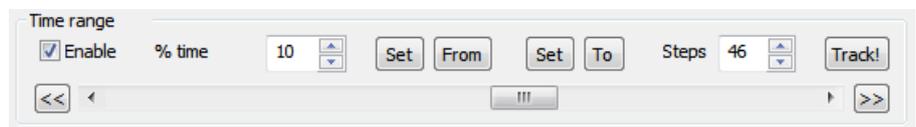


The basis idea is that each visible event is assigned to the cluster that it is most likely to belong to. The shape of each cluster is taken into account, so a cluster need not be spherical or aligned with the x, y or z axes. Probability is measured in terms of the Mahalanobis distance, which is the multi-dimensional equivalent of standard deviation.

The Match to classes within field sets the maximum separation in terms of the Mahalanobis distance between a cluster and an event, for the event to be considered as a possible cluster member. Events that do not belong to any cluster are given code 00. The matching does not pay any attention to the number of items in each class (unlike the Normal Mixtures from Existing command). The matching does take into account the fact that the class ellipsoids may not be aligned with the x, y and z axes (unlike the K Means from Existing command).

Time range control

This command is not available in online update mode. You can hide and show the Time range control area with the View menu Show Time Range command or with the `Ctrl+T` keyboard short-cut. This area restricts the number of displayed events by setting a time range that they must lie within. The % time range field sets the time range as a percentage of the time range spanned by all the events that were analysed. The thumb of the scroll bar is set to a size to match the time range.



When you check the Enable checkbox, the events displayed in the cluster window and used to build cluster statistics (and hence the class ellipses) are restricted in time. You can drag the thumb of the scroll bar to change the time range. Click the >> or << buttons the scroll bar will step through the data in the indicated direction; click again to stop it. You can use this to see if the clusters change position with time.

When you change the time range, everything that depends on the displayed events changes to track the displayed time range. If you have an active interval histogram and you have selected automatic updates in the INTH settings dialog, the histogram will update to track the current time range.

Tracking clusters

If you have clusters that move with time you can use the Track! button to assign class codes automatically over a time range. Proceed as follows:

1. Move the thumb to the start of the time range to track and make your initial cluster assignments, then click the Set next to the From button to mark the start position.
2. Move the thumb to the end of the time range (this can be before or after the From position time) and click the Set next to the To button to mark the end position.
3. Set the Steps field to the number of intermediate steps to use in progressing from the From position to the To position.

4. Click **Track!** to class events based on the clusters at the **From** position. The **K Means** dialog will open and you can choose the method to use for normalising the data. Once this is done, click **OK** to start tracking.

The **From** and **To** buttons jump the scroll bar to the start and end of the time range. The tracking algorithm works as follows:

1. The clusters at the **From** position are assumed to be correct. The statistics of these clusters are calculated. All the events past the start position to the end of the time range are marked as code **FF** – which stands for unknown.
2. The time range is moved on by: $(\text{To time} - \text{From time}) / \text{Steps}$ and we run the **K Means** algorithm on the new data and recalculate the statistics. The classes of events more than 3 times the Mahalanobis distance from the cluster centres are not changed.
3. Repeat step 2 until we reach the **To** time. Any events that are still marked as unknown are set to code **00**.

The larger the number of steps, the better the algorithm will cope with rapid cluster movement, but the longer the process will take. Because the algorithm is based on **K Means**, all the problems that **K Means** suffers from also apply. If you have periods where a class of events vanishes, or if new events appear in the middle of a time range, you may need to break down the tracking into multiple time ranges.

If there are clusters that do not move, you can greatly simplify tracking by coding these first (you may need to do this with the **Set Codes** command) and then hiding them.

Online clustering

When you open the clustering dialog while sampling data, you can choose to have new events added into the display as they occur (online update) or to run exactly as if you were analysing a file offline. You can enable and disable online update mode, set how far back in time to search for events, and limit the number of events to display with the **File** menu **Online Update** dialog.

If the **Edit WaveMark** dialog is in **At End** mode when the cluster dialog opens, the clustering starts in online update mode, and processes the most recent range of data as specified in the **Online Update** dialog. Otherwise, the range of data set for the **Edit WaveMark** dialog is processed. This initial analysis sets the parameters for online analysis. Changes in the **Edit WaveMark** dialog have no effect on online updates; you must use the **Reanalyse** command to pick up such changes.

Principal component analysis

Principal component analysis is complex and time-consuming. It is not practical to re-analyse all the data for each added event. Online updates use the matrices created by the initial analysis or by the last **Reanalyse** command to map new spikes onto the screen. The added spikes do not change the principal components.

Measurements

This method works identically online and offline.

Cluster on correlation and Cluster on errors

These methods take a snapshot of the templates in the **Edit WaveMark** dialog when the cluster windows are created and each time you **Reanalyse** the spikes. The online update mode uses these snapshots, not the templates in the **Edit WaveMark** dialog.

Getting started with clustering

We would suggest that you start by trying out principal component analysis, as this avoids you making subjective decisions about which features of your data are more important than others. It will also tend to produce clusters that are suitable for use with automatic clustering methods.

In the Edit WaveMark dialog choose the Analysis menu Principal Components command. This will open the cluster dialog and display the data, ready to cluster. You can check that visible clusters are stable over time by using the Time range control area. There are three basic ways to cluster the data:

Automatic	<p>If you have reasonably well-defined clusters that are, or can be made more or less spherical by scaling, and that contain similar numbers of events, then the K Means command will be able to separate them for you. This will often be the case with principal component values.</p> <p>If K Means does not make sensible choices, you should try the Normal Mixtures algorithm. This is slower, but can give better results.</p>
Manual	<p>If the data is chaotic, the only sensible course may be to manually position ellipses or user-defined shapes and set codes. You can position an ellipse on a cluster centre by right clicking at the centre and selecting a small, medium or large ellipse from the context menu.</p>
Semi-automatic	<p>If the automatic methods almost do what you want, you may be able to help them by marking the centres of clusters manually, then using one of the ...from existing commands to iterate from your centres.</p>

Whichever method you use, you may also wish only to mark events that you feel "certain" about. You can set events that are too far from the cluster centres to code 00 with the **Match to classes** command. Remember that changes made in the clustering dialog have no effect on the original data until you use the File menu **Apply** command.

K Means algorithm

The K Means algorithm has many variants, but the basic idea is based on the following iteration sequence.

1. Given the current set of K cluster centres; assign each event to the nearest centre.
2. Recalculate the K cluster centres based on the new assignments.
3. If the centres changed in the iteration, go back to 1, else done.

The algorithm always converges, but there is no guarantee that it converges to the optimum solution. The result depends on the original cluster centres. The **K Means for existing** command takes the current set of cluster centres as the starting point. The **K Means** command assigns events to clusters randomly before starting the algorithm and repeats the entire procedure 10 times. It chooses the solution that maximises the variance of the distance between all clustered points and the centre of gravity of all the points divided by the variance of the distance between each point and the centre of its cluster. The better the clusters are separated, the higher the value. This measure is known as J3 in some literature.

Once we have iterated to a solution the next task is to remove outlier points. We generate the statistics for each cluster and based on the assumption that the data follows a multivariate normal distribution, we set all points that lie more than 3 times the Mahalanobis distance away from the centre of their assigned cluster to have class code 00. We then run the K Means algorithm again, starting with the current centres. This usually converges in one or two iterations because removing outliers should make little difference to the distributions.

Limitations of KMeans

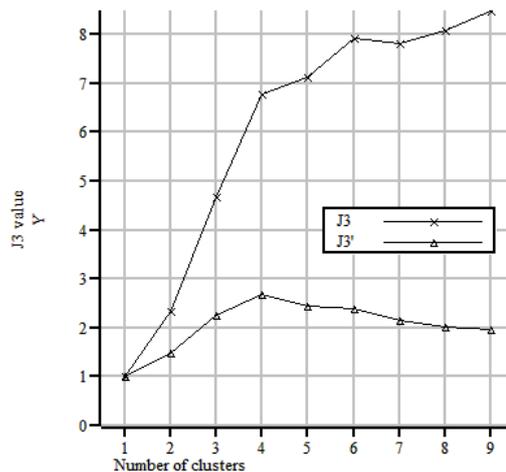
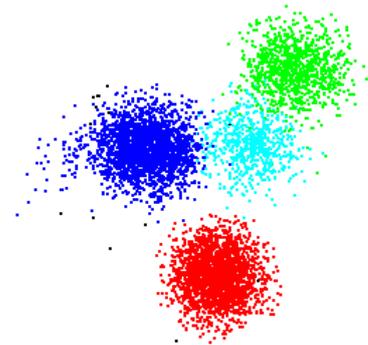
The KMeans algorithm assumes that the clusters are spherical and similar in size. The clustering is based on the distance between each point and the centre of each cluster. It is often the case, for example with clustering based on measurements, that the scaling of the axes are very different. You can compensate for this to some extent by normalising the data before clustering (shifting and scaling so that the data on each axis spans a similar range about zero). However, this will probably not make the clusters spherical.

Assuming that the clusters are more or less the same size and shape, they can be made spherical by stretching or shrinking the y and z axes relative to the x axis. This is done by applying weighting factors to the y and z axis data (the x axis data is assumed to have a weighting factor of 1.0). When data comes from the Principal component analysis, we guess the weights based on the results of the analysis. When data comes from measurements, the weights are set to 1. You can type in your own weights or weight the data based on how it looks on the display, or apply weights calculated so as to make the existing clusters as spherical as possible.

The KMeans algorithm also assumes that the x, y (and z) values are independent. If they are not, your data may look like ellipsoids that are not aligned with the x, y or z axis. If that is the case, you should probably use the Normal Mixtures algorithm for clustering.

J3 clustering measure

The J3 measure is displayed at the bottom of the clustering window after running the KMeans algorithm. It measures how compact the clusters are compared to the overall spread of the data. In this example there are 4 fairly clear clusters (although, at a stretch you could merge the two in the upper right corner together to make three clusters). The J3 values when analysing the data with different numbers of clusters are given in the following graph. The data was weighted using the existing cluster sizes, so we repeated the KMeans process until J3 stopped changing.



The J3 value is given by $J2/J1$ where J2 is the sum of the squares of the distances of all clustered points from the centre of all the points and J1 is the sum of the squares of the distances of each point from the centre of the cluster it belongs to. Dividing the data into more and more clusters reduces J1, so increases J3.

In this case, the J3 value increases sharply up to 4 clusters, then settles down at a fairly constant upward slope. Ideally we would like a measure of clustering that is a maximum at the most likely number of clusters.

We can estimate the value of J3 when data is divided up into more clusters than actually exist with the following simplistic argument. Consider the case where there are N points spread at a constant density through a m dimensional space. If this is divided up into k clusters of N/k points of similar (but arbitrary) shape, the radius (linear size) of each cluster will be proportional to $(N/k)^{1/m}$ and the variance of the distances from the centre of each cluster (J1) will be proportional to $(N/k)^{2/m}$. By the same argument, the J2 value will be proportional to $N^{2/m}$ with the same constant of proportionality, so J3 is given by $k^{2/m}$. So in the case where there is no way to divide data up into clusters (because it is spread through the space at a constant density), we know what the graph of J3 looks like.

In addition to J3, we also give the value J3', being J3 divided by $k^{2/m}$ which scales J3 by the value you would expect if all the points were uniformly distributed. In our case, we have $m=2$ (x and y axes) or $m=3$ (x, y and z axes). However, if one dimension makes no, or only a small contribution to the sorting, the effective number of dimensions is reduced so the J3' value can only be taken as a guide to the likely number of clusters.

Normal Mixtures algorithm

The Normal Mixtures algorithm assumes that the density of events around a cluster centre follows a multivariate normal distribution and that the number of clusters (K) is known. Clusters may be independent, or have the same shape, or have the same shape and be axially aligned or be spherical. We calculate the probability that an event belongs to a particular cluster based on the distances from each cluster, the shape of each cluster and the number of events in the cluster.

The algorithm is adapted from the Normal Mixtures algorithm described in *Clustering algorithms* by P.A. Hartigan, published in 1975 by Wiley, ISBN 0-471-35645-X. This process iterates towards a local maximum in the probability density of the set of events given the clusters. The value we return is the log of the probability density divided by the number of events. In outline, the algorithm proceeds as follows:

1. Make an initial assignment of the probabilities of each event belonging to each cluster. This can be based on the current classes, or can be done by randomly assigning events to clusters.
2. Using the probabilities of cluster membership, calculate the centres, shapes and number of events in each cluster.
3. Work out the probabilities of each event belonging to each cluster based on the new cluster centres and shapes.
4. Work out the most likely class code for each event. If any event is more than 3 times the Mahalanobis distance from the centre of the most likely class, set class 00.
5. If the most likely class assignment has not changed since the previous iteration or the iteration count is too high then stop, otherwise go back to step 2.

As set out, the algorithm converges slowly, so there are additional steps (not described here) required to accelerate the process. There is always the possibility that one or more classes will not be represented when iterating has finished. In this case, we take the class with the largest number of members and split this in two, then repeat this process until we have the desired number of classes, then we restart the iteration process.

There is no guarantee that any solution is the best as the outcome depends on the initial assignment of probabilities. If you choose to run the algorithm with no use of existing classifications, we repeat the process 10 times with randomised initial conditions and we choose the result that has the maximum probability density.

Mahalanobis distance

When describing clustering and the ellipses drawn to illustrate class boundaries we have referred to the Mahalanobis distance (named after P.C. Mahalanobis who described this measure in the 1930s). Given multivariate data values for which the values in each variable are normally distributed around a mean, this measure allows us to define boundaries of constant probability around the multi-dimensional centre of the distribution.

In one dimension, the normal distribution leads to a probability density $P(x)$ that is proportional to:

$$P(x) \propto \exp(-1/2x^2/v)$$

where x is the distance from the mean and v is the variance (σ^2). In two and three dimensions, this becomes:

$$P(x, y) \propto \exp(-1/2(x^2/v_{xx} + 2xy/v_{xy} + y^2/v_{yy}))$$

$$P(x, y, z) \propto \exp(-1/2(x^2/v_{xx} + 2xy/v_{xy} + y^2/v_{yy} + 2yz/v_{yz} + z^2/v_{zz} + 2zx/v_{zx}))$$

where x , y and z are the distances from the means and v_{ab} is the joint variance in the two components a and b . These equations generalise to any number of dimensions and are often written using matrix notation:

$$P(\mathbf{x}) \propto \exp(-\frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x})$$

where \mathbf{x} is the vector of distances from the mean, \mathbf{x}^T is the transpose of this, and \mathbf{C}^{-1} is the inverse of the covariance matrix.

The boundaries of constant probability for one, two and three dimensions satisfy the equations:

$$x^2 / \sigma^2 = r^2$$

$$x^2/v_{xx} + 2xy/v_{xy} + y^2/v_{yy} = r^2$$

$$x^2/v_{xx} + 2xy/v_{xy} + y^2/v_{yy} + 2yz/v_{yz} + z^2/v_{zz} + 2zx/v_{zx} = r^2$$

where r is a constant. By choosing suitable values of r we can set boundaries within which we would expect to find a certain proportion of the data. In one dimension, the boundary is a distance from the centre, in two dimensions it is an ellipse and in three dimensions this is an ellipsoid. When r is 1.0, we refer to the distance of constant probability density as the Mahalanobis distance. In the one-dimensional case, this is equivalent to the standard deviation.

Digital filtering

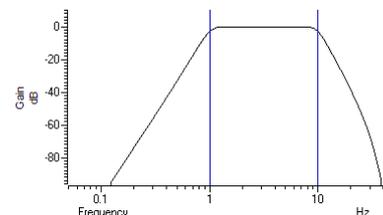
FIR and IIR filters

Filtering is used to remove unwanted frequency components from waveforms and can also be used to differentiate a signal. In Spike2 we provide you with two basic types of filter: Finite Impulse Response (FIR) and Infinite Impulse Response (IIR). Both types of filter have their advantages and disadvantages.

IIR filters

These are similar to analogue filters, and we design them by mapping standard Butterworth, Bessel, Chebyshev filters and resonators into their digital forms. IIR filters have these advantages:

- They can generate much steeper edges and narrower notches than FIR filters for the same computational effort.
- IIR filters are causal; they do not use future data to calculate the output, so there is no pre-ringing due to transients.



They also have disadvantages:

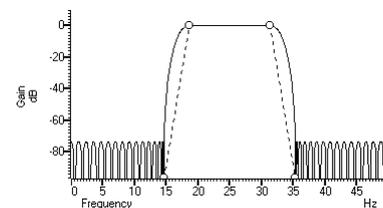
- IIR filters are prone to stability problems particularly as the filter order increases or when a filter feature becomes very narrow compared to the sample rate.
- IIR filters impose a group delay on the data that varies with frequency. This means that they do not preserve the shape of a waveform, in particular, the positions of peaks and troughs will change.

The output of an IIR filter may take a long time to settle down from the discontinuity at the start (transition from no data to the supplied data).

FIR filters

We describe FIR filters in terms of frequency bands: *pass bands*, *stop bands* and *transition gaps*. You define a filter by the arrangement of bands and the corner frequencies of each band. FIR filters have these advantages:

- They are unconditionally stable as they do not feedback the output to the input.
- There is no phase delay through the filter, so peaks and troughs do not move when data is filtered (this is called *linear phase* in the literature).



They also have disadvantages:

- They are poor at generating very narrow notches or narrow band pass filters.
- The narrowest frequency band or band gap is limited by the number of coefficients (we allow up to 511 coefficients).
- FIR filters are not causal; they use future as well as past data to generate each output point. A transient in the input causes effects in the output before the transient.

If your FIR filter has n coefficients, the first and last $n/2$ output points are estimates due to the discontinuity at the start and end of the data.

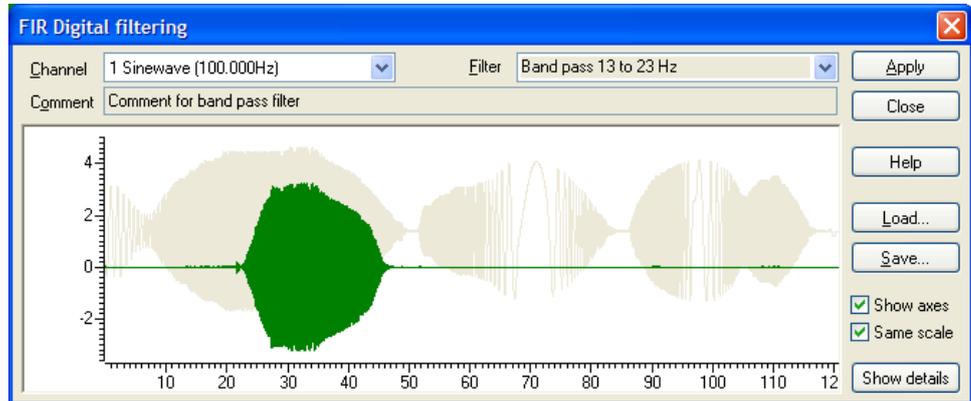
So which to choose?

If you want a differentiating filter, you have no choice as we have not implemented differentiators for IIR filters. Unless one of the disadvantages of the FIR filter is a problem, you will likely have fewer unexpected effects with an FIR filter.

However, there are circumstances in which only an IIR filter will do. If you need a high Q notch filter or resonator, then use an IIR filter. If you are interested in small changes just before a large discontinuity, only the IIR filter will help you. However, make sure that you understand the disadvantages of IIR filters before you depend on their output.

Digital filter dialog

The Analysis menu Digital filters command is available when you have a data file open that contains waveform channels. You can choose to apply Finite Impulse Response (FIR) filters or Infinite Impulse Response (IIR) filters. Apart from the dialog title, the initial dialog display is the same for IIR and FIR filters. You can apply one of twelve stored digital filters to a waveform channel, or you can create your own digital filter. You can also load and save additional sets of filters from the dialog.



The dialog shows the original waveform in grey, and a filtered version in the colour you have set for waveform data in a time view. Whenever you change the filter, the display updates to show the effect of the change.

Show axes You can choose to **Show axes** for the original data and for the filtered version. The axis for the original data is always on the left. If a separate axis is required for the filtered data, it is drawn in the filtered data colour on the right.

Same scale Initially, the filtered data is drawn at the same scale as the original data. However, sometimes this is inconvenient, for example when high-pass filtering a signal with a significant DC offset or when the result of filtering is very small compared to the original. If you clear the **Same scale** checkbox, the filtered data is scaled to fit in the window independently of the scaling of the original waveform.

Channel The Channel field allows you to select a Waveform or a RealWave channel to filter.

Filter The Filter field of the dialog box selects the filter to apply. There are normally 12 filters to choose from. When you first open the dialog, this field is grey, indicating that you cannot edit the filter name. If you display the filter details you can modify the filter name.

Comment The Comment field is for any purpose you wish; there is one comment per filter. When you first open the dialog, this field is grey, indicating that you cannot edit the comment. Click the **Show details** button to edit the comment.

The Filter field of the dialog box selects the filter to apply and the Channel field sets the waveform channel to filter.

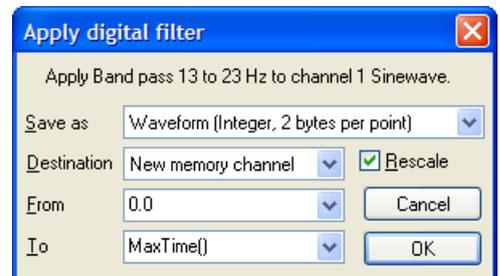
The Close button shuts the dialog and will ask if you want to save any changed filter and the Help button opens the on-line Help at the digital filtering topic.

Close Click the **Close** button to shut the filtering dialog. If you have made a change to any of the filters, or loaded a new filter set, you are asked if you want to save the current set of filters as your standard filter set.

Load and Save You can choose to save the current set of filters to a `.cfb` file, or to load a new set of filters from a `.cfb` file. If you are working with FIR filters, this only saves or loads FIR filters. If you are working with IIR filters, this only saves or loads IIR filters. Loading a new filter set does not change your standard filter set, however, you will be asked if you want to change your standard set when you close the dialog.

Show details The Show details button increases the dialog size to display a new area in which you can design and edit filters. Click this button again to hide the new dialog area. When you display the filter details, the Filter and Comment fields become editable. If you change a filter or create a new filter or load a new set of filters from a file, you will be prompted to save the filter bank when you close the digital filter dialog. The details are different for FIR and IIR filters.

Apply The Apply button opens a new dialog in which you set where to write the result of the filter operation on the data channel. In the IIR filter dialog, this button is disabled if the filter is not valid. The channel comment of the new channel holds the source channel number and a description of the filtering operation.



Save As You can save the results of filtering as 16-bit integer data (Waveform) using a scale and offset to convert to user units, or as 32-bit floating point data (RealWave). RealWave output gives you a more accurate result at the cost of using twice as much storage space.

Destination You can write the result to a new memory channel, to an unused channel on disk, or to a memory channel that holds waveform or RealWave data with the same sampling rate as the channel you are filtering.

From and To The From and To fields set the time range to process.

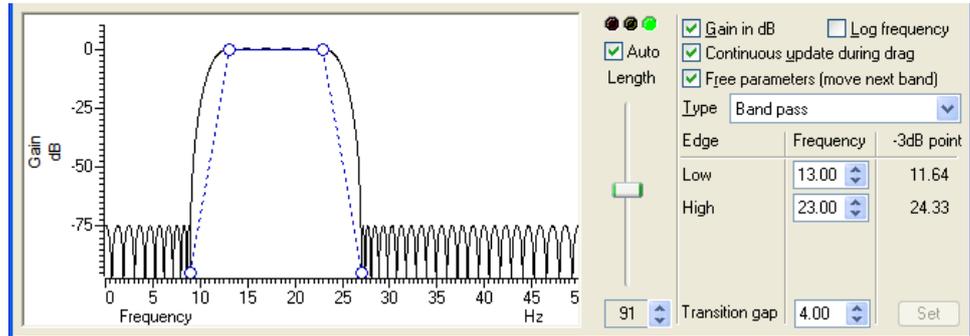
Rescale If you check Rescale, the output data scale and offset are set to give the best possible representation of the waveform as 16-bit integers. If you save the output as a Waveform, this doubles the time required for filtering. If you Rescale when adding data to an existing memory channel, the scale and offset take into account both the added data and any remaining original data. If Rescale is unchecked, the scale and offset values from the source channel are copied. You would normally wish to rescale output to a RealWave channel as this improves accuracy if the channel is ever read as integer data.

OK Click OK to start filtering. As applying a filter can be a lengthy process, a progress dialog appears with a Cancel button during the filtering operation.

Filter bank A digital filter definition is complex and it would be tedious to specify all the properties of a filter each time you wanted to apply one to data. To avoid this, Spike2 contains a filter bank of 12 FIR and 12 IIR filter definitions. This filter bank is saved to the file `Filtbank.cfb` when you close Spike2 and reloaded when you open it. When you use the digital filter dialog, you specify which filter you want by the filter name.

Script users identify the filter by its type (FIR or IIR) and an index number in the range 0 to 11. Script users also have access to two additional temporary filters with index number -1 (one for FIR and the other for IIR filters) that they can set and use for channel filtering operations without changing the standard filter set.

FIR filter details



The graph in the details area displays the ideal and actual frequency response of the filter. The ideal response appears as blue solid lines for each pass band linked by dotted lines that mark transition gaps between the bands. All transition gaps have the same width. The calculated frequency response is drawn as solid black lines and is greyed when the filter specification has changed and the response has not yet been calculated.

The mouse pointer changes to indicate the feature it is over: $\leftarrow\rightarrow$ for a pass band or stop band, $\leftarrow\circ\rightarrow$ for a transition gap and $\leftarrow\circ$ for a band corner. The small circles can be dragged sideways to change the slope of the band edges or you can edit the band edges as numbers in the Frequency panel on the right. You can also drag the bands and the transition gaps sideways.

Set If you edit the numbers in the Frequency panel, the Set button is enabled so you can force a recalculation of the filter.

-3dB point The filters produced by the program are not defined in terms of -3dB corner frequencies and n dB per octave, as is often the case for traditional analogue filters. The -3dB point column is present to help users who are more comfortable describing filter band edges in terms of the 3 dB point.

Gain in dB The Gain in dB check box sets the y axis scale to dB if checked, linear if not checked. Using dB is usually the more convenient, except when working on a differentiator.

Log frequency You can choose to display the frequency axis as the logarithm of the frequency, which gives you more resolution at low frequencies. However, this can make working with FIR filters more awkward as it removes the visual symmetry of the transition bands. In log mode, the frequency axis extends down to 0.001 of the data sample rate.

Continuous update If you check the Continuous update box, the filter is updated while you drag the filter features around. If you have a slow computer and this feels ponderous you can clear the check box, in which case the filter is not recalculated until you stop changing features.

Free parameters If you check the Free parameters box, dragged features are not limited by the next band and will push bands along horizontally. If you clear the box, the horizontal motion of a dragged feature is limited by the next moveable object.

Length, Auto and traffic lights To the right of the frequency response display is a slider that controls the number of filter coefficients. In general, the more coefficients, the better the filter. However, the more coefficients, the longer it takes to compute them and the longer to filter the data. If you check the Auto box, the program will adjust the number of coefficients for you to produce a useful filter (with around 70 dB cut between pass and stop bands). The “traffic light” display above the slider shows green if the filter is good, amber if the result is usable but not ideal, and red if the result is hopeless.

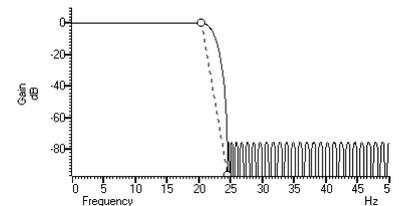
An FIR filter of length n uses the $n/2$ points before and after each input point to produce each output point. When there is no input data available before or after an input point, the filter uses a duplicate of the nearest point as an estimate of the data value. The $n/2$ output points next to any break in the input data should not be used for any critical purpose.

FIR Filter types The **Type** field sets the arrangement of filter bands. If you need a filter that is not in this list you can generate it from the script language with the `FIRMake()` command. There are currently 12 different filter types:

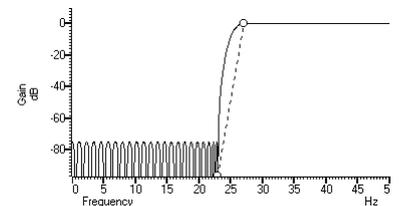
All pass This has no effect on your signal. This filter type covers the case where you apply a low pass filter designed for a higher sampling rate to a waveform with a much lower sampling rate, so that the pass band extends beyond half the sampling frequency of the new file.

All stop This removes any signal; the output is always zero. This filter type is provided to cover the case where you apply a high pass filter designed for a higher sampling rate to a waveform with a much lower sampling rate, so that the stop band extends beyond half the sampling frequency of the new file.

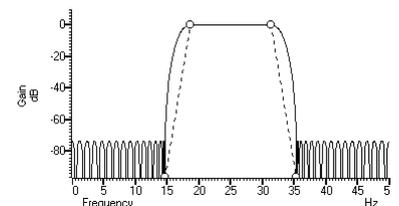
Low pass This filter attempts to remove the high frequencies from the input signal. The **Frequency** field holds one editable number, **Low pass**, the frequency of the upper edge of the pass band. The stop band starts at this frequency plus the value set by the **Transition gap** field.



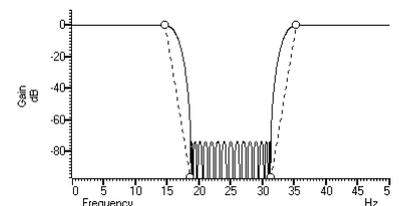
High pass A high pass filter removes low frequencies from the input signal. The **Frequency** field holds one editable number, **High pass**, the frequency of the lower edge of the pass band. The stop band starts at this frequency less the value set by the **Transition gap** field.



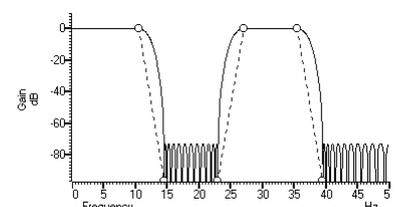
Band pass A band pass filter passes a range of frequencies and removes frequencies above and below this range. The **Frequency** field has two editable numbers, **Low** and **High**, which correspond to the two edges of the pass band. The stop band below runs up to **Low-Transition gap**, and the stop band above from **High+Transition gap** to one half the sampling rate.



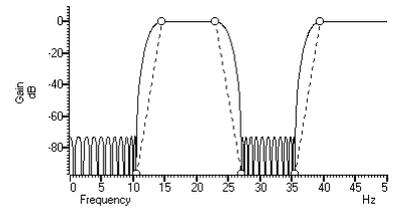
Band stop A band stop filter removes a range of frequencies. The **Frequency** field has two editable numbers, **High** (the upper edge of the first pass band) and **Low** (the lower edge of the upper pass band). The stop band below runs from **High+Transition gap** up to **Low-Transition gap**.



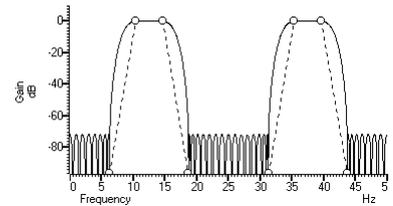
One and a half low pass This filter has two pass bands, the first running from zero Hz and the second in the frequency space between the upper edge of the first pass band and one half the sampling rate. The **Frequency** field has three editable numbers: **Band 1 high**, **Band 2 low** and **Band 2 high**. These numbers correspond to the edges of the pass bands.



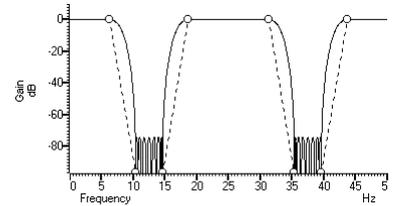
One and a half high pass This filter has two pass bands. The second runs up to one half the sampling rate. The first band lies in the frequency space between 0 Hz and the lower edge of the second band. The Frequency field has three editable numbers: Band 1 low, Band 1 high and Band 2 low. These numbers correspond to the edges of the pass bands.



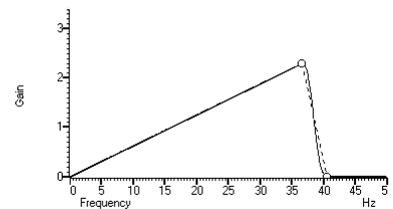
Two band pass This filter passes two frequency ranges and rejects the remainder. Both 0 Hz and one half the sampling frequency are rejected. The Frequency field has 4 numeric fields: Band 1 low, Band 1 high, Band 2 low and Band 2 high. These fields correspond to the four edges of the two bands.



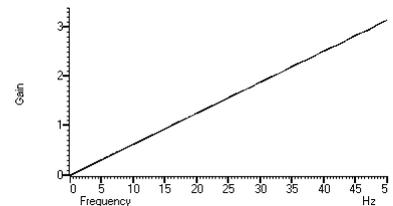
Two band stop This filter passes three frequency ranges and rejects the remainder. Both 0 Hz and one half the sampling rate are passed. The Frequency field has 4 numeric fields: Band 1 high, Band 2 low, Band 2 high and Band 3 low. These fields correspond to the four edges of the three bands.



Low pass differentiator This filter is a combination of a differentiator (that is the output is proportional to the rate of change of the input) and a low pass filter. The y axis scale is linear, rather than in dB (although you can display it in dB if you wish). There is one editable number in the Frequency field, Low pass, the end of the differential section of the filter.



Differentiator The output of the filter is proportional to the rate of change of the input. The y axis scale is linear, rather than in dB (although you can display it in dB if you wish). The Frequency field is empty as there is only one band and it extends from 0 Hz to half the sampling rate.



IIR filter details

We describe IIR filters in terms of a filter type (low pass, high pass, band pass or band stop), the analogue filter model that they are based on (Butterworth, for example), the corner frequencies and the filter order (which determines the steepness of the cut-off outside the desired pass bands). Filters based on Chebyshev designs also require a ripple specification and resonators require a Q factor.

The graph in the details area displays the corner frequencies and the frequency response of the filter. You can adjust the filter by clicking and dragging in this area or by editing the filter parameters as text. The mouse pointer changes to indicate the feature it is over:  for a corner frequency and  for an adjustable parameter (ripple or Q factor).

Gain in dB The **Gain in dB** check box sets the y axis scale to dB if checked, linear if not checked. Using dB is usually the more convenient.

Log frequency You can choose to display the frequency axis as the logarithm of the frequency, which gives you more resolution at low frequencies. The display extends to 0.0001 of the sample rate. If you need a corner frequency below this you should consider sampling the signal more slowly in the first place. Alternatively, low pass filter the signal and then use a channel process to down sample it before filtering.

Continuous update This is always checked for IIR filters, and is greyed out.

Free parameters If you check the **Free parameters** box, corner frequencies are not limited by the next corner, and will push them along. If you clear the box, corner frequencies cannot be moved past each other.

Traffic lights and messages The traffic lights show green if the filter appears to be OK, amber if the filter may be unstable and red if the filter calculation failed, the filter is unstable or if one of the input parameters is illegal. There will usually be an explanatory message in the lower right hand corner of the dialog explaining the problem. In the amber state, the filter may still be usable; you can look at the frequency response and the filtered data to see if the result is acceptable.

Filter Order In an analogue filter, the filter order determines the sharpness of the filter, for example Butterworth and Bessel filters tend towards $6n$ dB per octave, where n is the filter order. In a digital filter, the order determines the number of filter coefficients. The higher the order, the sharper the filter cut-off and also, the more likely the filter is to be unstable due to problems in numerical precision. You can set filter orders of 1 to 10. You should always use the lowest order that meets your filtering criteria. Phase non-linearity gets worse as the filter order increases.

Filter type There are four filter types: **Low pass**, **High pass**, **Band pass** and **Band stop**. However, if you set the filter model to **Resonator**, there are only two types, being **Band pass** (this creates a resonator filter) and **Band stop** (this creates a notch filter). For all filter models except Chebyshev type 2 and Resonator, the frequencies given are the points at which the filter achieves a cut of 3 dB. For Chebyshev type 2 filters, the frequencies are the point at which the filter cut reaches the value set by the Ripple parameter. For Resonators, the frequency is the centre frequency of the resonator.

Low pass Use this to remove high frequencies and pass low frequencies. This has a single corner frequency.

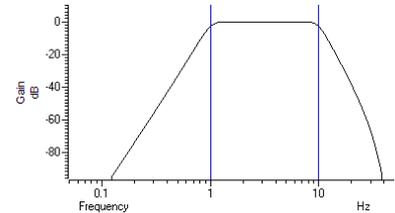
High pass Use this to remove low frequencies and pass high frequencies. This has a single corner frequency.

Band pass This passes a range of frequencies between the Low edge and the high edge. If the filter model is Resonator, then this has a single Centre frequency.

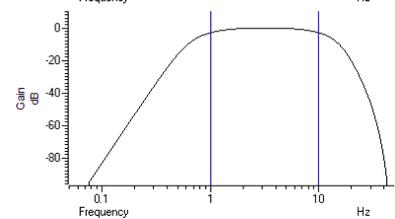
Band stop This removes a range of frequencies between the Low edge and the high edge. If the filter model is Resonator, then this has a single Centre frequency.

Filter Model The IIR filters we provide are digital implementation of standard analogue filter models. The following descriptions use fifth order 1 to 10 Hz band pass filters on 100 Hz data as examples (except for the Resonator filters). The five filter models are:

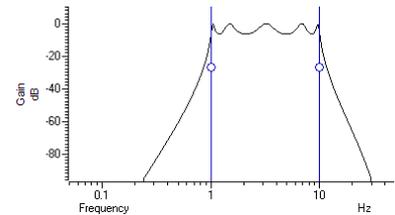
Butterworth This has a maximally flat pass band, but pays for it by not having the steepest possible transition between the pass band and the stop band. This is a good choice for a general-purpose IIR filter, but beware that the group delay can get quite bad near the corners, especially for high-order filters.



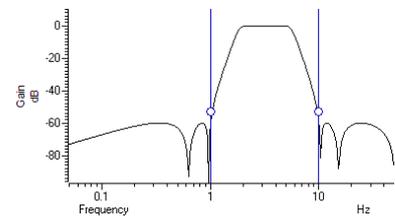
Bessel An analogue Bessel filter has the property that the group delay is maximally flat, which means that it tends to preserve the shape of a signal passed through it. This leads to filters with a gentle cut-off. When digitised, the constant group delay property is compromised; the higher the filter order, the worse the group delay.



Chebyshev type 1 Filters of this type are based on Chebyshev polynomials and have the fastest transition between the pass band and the stop band for a given ripple in the pass band and no ripples in the stop band. You can adjust the ripple by dragging the small circles vertically or with the Ripple field to the right of the displayed frequency response.

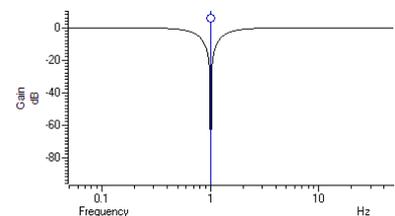


Chebyshev type 2 Filters of this type are defined by the start of the stop band and the stop band ripple (the minimum cut in the stop band). They have the fastest transition between the pass and stop bands given the stop band ripple and no ripple in the pass band. Drag the small circles to adjust the ripple, or use the Ripple field to the right of the frequency response.

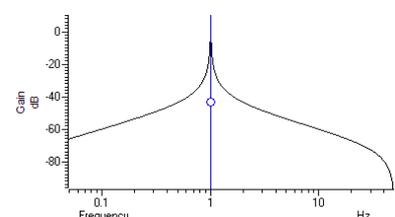


Resonator Resonators are defined by a centre frequency and a Q factor. The Q is the centre frequency divided by the width of the resonator at the -3 dB point. You can have a band stop or a band pass resonator. The Q is adjusted by dragging the small circle or with the Q field to the right of the displayed frequency response.

Band stop resonators are also called Notch filters. The higher the Q, the narrower the notch. Notch filters are sometimes used to remove mains hum, but if you do this you will likely need to set notches at the first few odd harmonics of the mains frequency; notch filters are not recommended if you care about the phase of signals in the region of the notch centre frequency. The example has a very low Q (1.24) to make the filter response visible.

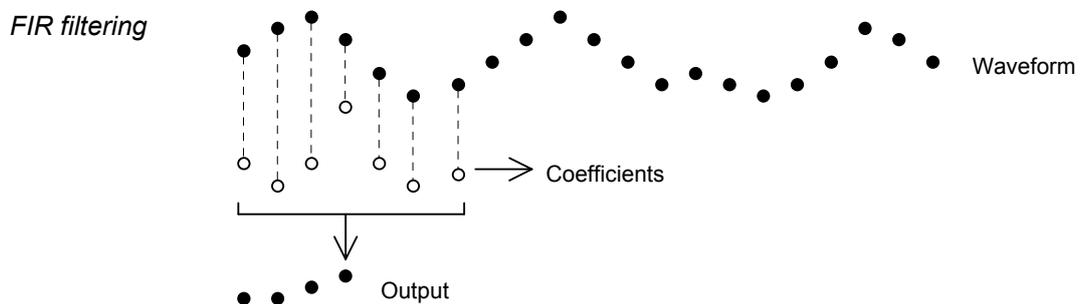


A band pass resonator is the inverse of a notch. Band pass resonators are sometimes used as alternatives to a narrow bandpass filter. The example has a Q of 100. The higher the Q set for a resonator, the longer it will take for the output to stabilise at the start of the filter output.



FIR filters The `FIRMake()`, `FIRQuick()` and `FiltCalc()` script commands and the **Analysis** menu **Digital filters...** dialog generate FIR (Finite Impulse Response) filter coefficients suitable for a variety of filtering applications. The generated filters are optimal in the sense that they have the minimum ripple in each defined band. These filter coefficients are used to modify a sampled waveform, usually to remove unwanted frequency components. The algorithmic heart of the filter coefficient generation is based on the well-known FORTRAN program written by Jim McClellan of Rice University in 1973 that implements the *Remez exchange algorithm* to optimise the filter.

The theory of FIR filters is beyond the scope of this document. Readers who are interested in learning more about the subject should consult a suitable text book, for example *Theory and Application of Digital Signal Processing* by Rabiner and Gold published by Prentice-Hall, ISBN 0-13-914101.



This diagram shows the general principle of the FIR filter. The hollow circles represent the filter coefficients, and the solid circles are the input and output waveforms. Each output point is generated by multiplying the waveform by the coefficients and summing the result. The coefficients are then moved one step to the right and the process repeats.

From this description, you can see that the filter coefficients (from right to left) are the *impulse response* of the filter. The impulse response is the output of a filter when the input signal is all zero except for one sample of unit amplitude. In the example above with 7 coefficients, there is no time shift caused by the filter. With an even number of coefficients, there is a time shift in the output of half a sample period.

Frequencies The **Analysis** menu **Digital filters...** command deals with frequencies in Hz as this is comfortable for us to work with. However, if you calculate a FIR filter for one sampling rate, and apply the same coefficients to a waveform sampled at another rate, all the frequency properties of the filter are scaled by the relative sampling rates. That is, the frequency properties of an FIR filter are invariant when expressed as fractions of the sampling rate, not when expressed in Hz.

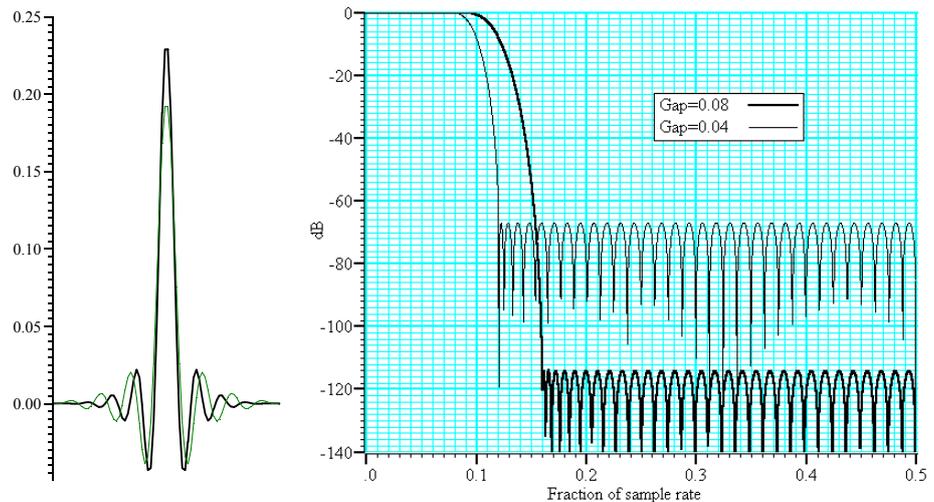
It is usually more convenient when dealing with real signals to describe filters in terms of Hz, but this means that each time a filter is applied to a waveform the sampling rate must be checked. If the rate is different from the rate for which the filter was last used, the coefficients must be recalculated. Unless you use the `FIRMake()` script command, Spike2 takes care of all the frequency scaling and recalculation for you. The remainder of this description is to help users of the `FIRMake()` script command, but the general principles apply to all the digital filtering commands in Spike2.

Users of the `FIRMake()` script command must specify frequencies in terms of fractions of the sample rate from 0 to 0.5. For example, if you were sampling at 10 kHz and you wanted to refer to a frequency of 500 Hz, you would call this 500/10000 or 0.05.

Example filter The heavy lines in the next diagrams show the results obtained by `FIRMake()` when it designed a low pass filter with 80 coefficients with the specification that the frequency band from 0 to 0.08 should have no attenuation, and that the band from 0.16 to 0.5 should be removed. We can specify the relative weight to give to the ripple in each band. In this case, we said that it was 10 times more important that the *stop band* (0.16 to 0.5) should pass no signal than the *pass band* should be completely flat.

We have shown the coefficients as a waveform for interest as well as the frequency response of the filter. The shape shown below is typical for a band pass filter. One way of understanding the action of the FIR filter is to think of the output as the correlation of the waveform and the filter coefficients.

Coefficients and frequency response for low pass filters



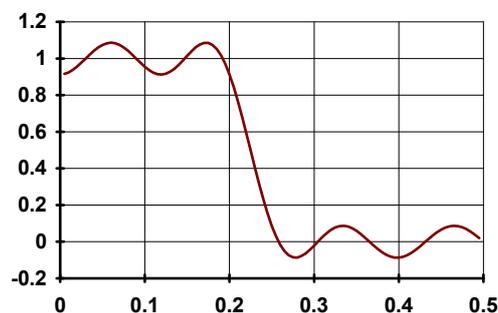
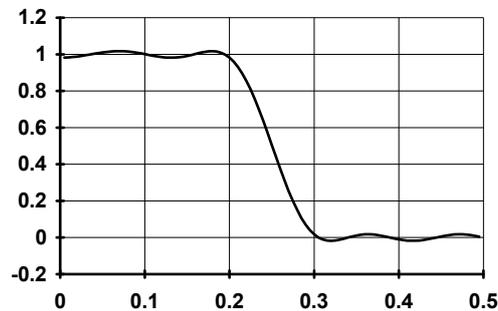
The frequency response is shown in dB, which is a logarithmic scale. A ratio r is represented by $20 \log_{10}(r)$ dB. A change of 20 dB is a factor of 10 in amplitude, 6 dB is approximately a factor of 2 in amplitude. The graph shows that a frequency in the stop band is attenuated by over 110 dB (a factor of 300,000 in amplitude with respect to the signal before it was filtered).

Because we didn't specify what happened between a frequency of 0.08 and 0.16 of the sampling rate, the optimisation pays no attention to this region. You might ask what happens if we make this transition gap smaller. The lighter line in the graph shows the result of halving the width of the gap by making the stop band run from 0.12 to 0.5. The filter is now much sharper. However, you don't get something for nothing. The attenuation in the stop band is reduced from 110 dB to around 70 dB. Although you cannot see it from the graph, the ripple in the pass band also increases by the same proportion (from 1 part in 30,000 to 1 part in 300).

We can restore the attenuation in the stop band by increasing the number of coefficients to around 120. However, there are limits to the number of coefficients it is worth having (apart from increasing the time it takes to calculate the filter and filter the data). Although the process used to calculate coefficients uses double precision floating point numbers, there are rounding errors and the larger the number of coefficients, the larger the numerical noise in the result.

Because the waveform channels are stored in 16-bit integers, there is no point designing filters that attenuate any more than 96 dB as this is a factor of 32768 (2^{15}). Attenuations greater than this would reduce any input to less than 1 bit. If you are targeting data stored in real numbers this restriction may not apply.

It is important that you leave gaps between your bands. The smaller the gap, the larger the ripple in the bands.



This is illustrated by these two graphs. They show the linear frequency response of two low pass filters, both designed with 18 coefficients (we have used so few coefficients so the ripple is obvious). Both have a pass band of 0 to 0.2, but the first has a gap between the pass band and the stop band of 0.1 and the second has a gap of 0.05. We have also given equal weighting to both the pass and the stop bands, so you can see that the ripple around the desired value is the same for each band.

As you can see, halving the gap has made a considerable increase in the ripple in both the pass band and the stop band. In the first case, the ripple is 1.76%, in the second it is 8.7%. Halving the transition region width increased the ripple by a factor of 5.

In case you were worrying about the negative amplitudes in the graphs, a negative amplitude means that a sine-wave input at that frequency would be inverted by the filter. The graphs with dB axes consider only the magnitude of the signals, not the sign.

FIRMake() filter types

`FIRMake()` can generate coefficients for four types of filter: Multiband, Differentiators, Hilbert transformer and a variation on multiband with 3 dB per octave frequency cut. The other routines can generate only Multiband filters and Differentiators.

Multiband filters

The filter required is defined in terms of frequency bands and the desired frequency response in each band (usually 1.0 or 0.0). Bands with a response of 1.0 are called *pass bands*, bands with a response of 0.0 are called *stop bands*. You can also set bands with intermediate responses, but this is unusual. The bands may not overlap, and there are gaps between the defined bands where the frequency response is undefined. You give a weighting to each band to specify how important it is that the band meets the specification. As a rule of thumb, you should make the weight in stop bands about ten times the weight in pass bands.

`FIRMake()` optimises the filter by making the ripple in each band times the weight for the band the same. The ripple is the maximum error between the desired and actual filter response in a band. The ripple is usually expressed in dB relative to the unfiltered signal. Thus the ripple in a stop band is the minimum attenuation found in that band. The ripple in a pass band is the variation of the frequency response from the desired response of unity. In some situations, for example audio filters, quite large ripples in the pass band are tolerable but the same ripple would be unacceptable in a stop band. For example, a ripple of -40 dB in a pass band (1%) is inaudible, but the same ripple in a stop band would allow easily audible signals to pass. By weighting bands you can increase the attenuation in one band at the expense of another to suit your application.

Differentiators The output of a differentiator increases linearly with frequency and is zero at a frequency of 0. The differentiator is defined in terms of a frequency band and a slope. The frequency response at frequency f is $f * slope$. The slope is usually set so that the frequency response at the highest frequency is no more than 1.

The weight given to each frequency within a band is the weight for that band divided by the frequency. This gives a more accurate frequency response at low frequencies where the resultant amplitude will be the smallest.

Although you can define multiple bands for a differentiator, it is unusual to do so. Almost all differentiators define a single band that starts at 0. Occasionally a differentiator followed by a stop band is needed.

Hilbert transformers A Hilbert transformer is a very specialised form of filter that causes a phase shift of $-\pi/2$ in a band, often used to separate a signal from a carrier. The theory and use of this form of filter is way beyond the scope of this document. Unless you know that you need this filter type you can ignore it.

Multiband with 3dB/octave cut This is a variation on the multiband filter that can be used to filter white noise to produce band limited pink noise. The filter is identical to the band pass filter except that the attenuation increases by 3 dB per octave in the band (each doubling of frequency reduces the amplitude of the signal by a factor of the square root of 2). It is used in exactly the same way as the multiband filter.

Low pass filter example A waveform is sampled at 1 kHz and we are interested only in frequencies below 100 Hz. We would like all frequencies above 150 Hz attenuated by at least 70 dB.

A low pass filter has two bands. The first band starts at 0 and ends at 100 Hz, the second band starts at 150 Hz and ends at half the sampling rate. Translated into fractions of the sampling rate, the two bands are 0-0.1 and 0.15 to 0.5. The first band has a gain of 1, the second band has a gain of 0. We will follow our own advice and give the stop band a weight of 10 and the pass band a weight of 1. We will try 40 coefficients to start with, so a possible script is:

```
var prm[5][2];      'Array for parameters
var coef[40];      'Array for the coefficients
'   band start      band end      function      weight
prm[0][0]:=0.00;  prm[1][0]:=0.1;  prm[2][0]:=1.0;  prm[3][0]:= 1.0;
prm[0][1]:=0.15;  prm[1][1]:=0.5;  prm[2][1]:=0.0;  prm[3][1]:=10.0;
FIRMake(1, prm[0][0], coef[0]);
PrintLog("Pass Band ripple=%.1fdB Stop band attenuation=%.1f\n",
        prm[4][0], prm[4][1]);
```

If you run this, the log view output is:

```
Pass Band ripple=-28.8dB Stop band attenuation=-48.8
```

The attenuation in the stop band is only 48 dB, which is not enough. The ripple in the pass band is around 3% of the signal amplitude. We can increase the stop band attenuation in three ways: by increasing the number of coefficients, by giving the stop band more weight, or by making the gap larger between the bands.

We don't want to give the stop band more weight; this would increase the ripple in the pass band. We could probably reduce the width of the pass band a little as the attenuation of the signal tends to start slowly, but we will leave that adjustment to the end. The best way to improve the filter is to increase the number of coefficients. If we increase the size of `coef[]` to 80 coefficients and run again, the output now is:

```
Pass Band ripple=-58.7dB Stop band attenuation=-78.7
```

This is much closer to the filter we wanted. You might wonder if there is a formula that can predict the number of coefficients based on the filter specification. There is no exact relationship, but the following formula, worked out empirically by curve fitting, predicts the number of coefficients required to generate a filter with equal weighting in each of the bands and is usually accurate to within a couple of coefficients. The formula can be applied when there are more than two bands, but becomes less accurate as the number of bands increase.

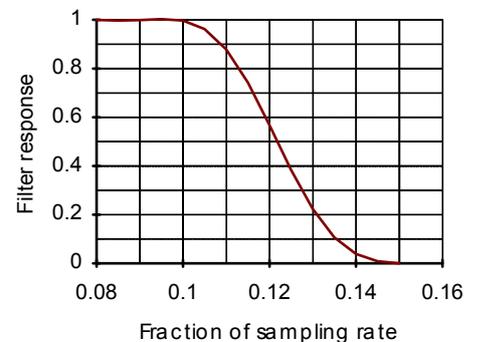
```
' dB      is the mean ripple/attenuation in dB of the bands
' deltaF is the width of the transition region between the bands
' return An estimate of the number of coefficients
Func NCoefMultiBand(dB, deltaF)
return (dB-23.9*deltaF-5.585)/(14.41*deltaF+0.0723);
end;
```

In our example we wanted at least 70 dB attenuation, and we weighted the stop band by a factor of 10 (20 dB). This causes a 10 dB improvement in the stop band at the expense of a 10 dB degradation of the pass band. Thus to achieve 70 dB in the stop band with the weighting, we need 60 dB without it. If we set these values in the formula ($dB = 60$, $\delta F = 0.05$), it predicts that 67.13 coefficients are needed. If we run our script with 67 coefficients, we get 70.9 dB attenuation, which is close enough!

A final finesse

If we look at the frequency response of our filter in the area between the pass band and the stop band, we see that the curve is quite gentle to start with. If you are used to using analogue filters, you will recall that the corner frequency for a low pass analogue filter is usually stated to be the frequency at which the filter response fell by 3 dB which is a factor of $\sqrt{2}$ in amplitude (when the response falls to 0.707 of the unfiltered amplitude).

If we use the analogue filter definition of corner frequency, we see that we have produced a filter that passes from 0 to 0.115 of the sampling rate, and we wanted from 0 to 0.1, so we can move the corner frequency back. This will increase the attenuation in the stop band, and reduce the filter ripple, as it widens the gap between the pass band and the stop band. If we move it back to 0.085, the attenuation in the stop band increases to 84 dB. Alternatively, we could move both edges back, keeping the width of the gap constant. This leaves the stop band attenuation more or less unchanged, but means that the start of the stop band is moved lower in frequency.

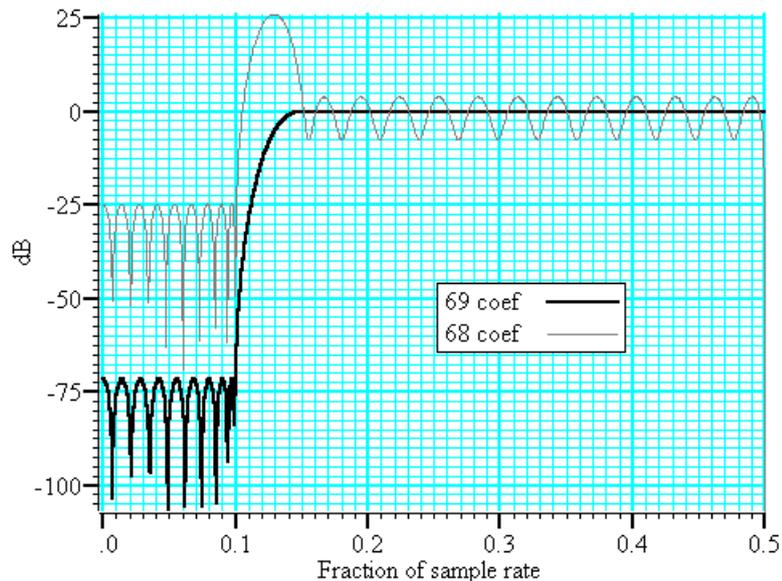


High pass filter

A high pass filter is the same idea as a low pass except that the first frequency band is a stop band and the second band is a pass band. All the discussion for a low pass filter applies, with the addition that **there must be an odd number of coefficients**. If you try to use an even number your filter will be very poor indeed. The example below shows a script for a high pass filter with the same bands and tolerances as for the low pass filter. We have added a little more code to draw the frequency response in a result view.

```
var prm[5][2];
var coef[69];
'   band start      band end      function      weight
prm[0][0]:=0.00; prm[1][0]:=0.1; prm[2][0]:=0.0; prm[3][0]:=10.0;
prm[0][1]:=0.15; prm[1][1]:=0.5; prm[2][1]:=1.0; prm[3][1]:= 1.0;
FIRMake(1, prm[0][], coef[]);
const bins% := 1000;
var fr[bins%];
FIRResponse(fr[], coef[], 0);
SetResult(bins%, 0.5/(bins%-1), 0, "Fr Resp", "Fr","dB");
ArrConst([], fr[]);
Optimise(0);
WindowVisible(1);
```

Effect of odd and even coefficients



The graph shows the results of this high pass filter design with 69 coefficients, which gives a good result, and with 68 coefficients, which does not. In fact, if we had not given a factor of 10 weight (20 dB) to the stop band, the filter with 68 coefficients would not have achieved any cut in the stop band at all!

The reason for this unexpected result is that we have specified a non-zero response at the Nyquist frequency (half the sampling rate). If you imagine a sine wave with a frequency of half the sample rate, each cycle will contribute two samples. The samples will be 180° out of phase, so if one sample has amplitude a , the next will have amplitude $-a$, the next a and so on. The filter coefficients are mirror symmetrical about the centre point for a band pass filter, so with an even number of coefficients, the result when the input waveform is $a, -a, a, -a, \dots$ is 0. Another way of looking at this is to consider that a filter with an even number of coefficients produces half a sample delay. The output halfway between points that are alternately $+a$ and $-a$ must be 0.

You can use the formula given for the low pass filter to estimate the number of coefficients required, but you must round the result up to the next odd number.

General multiband filter

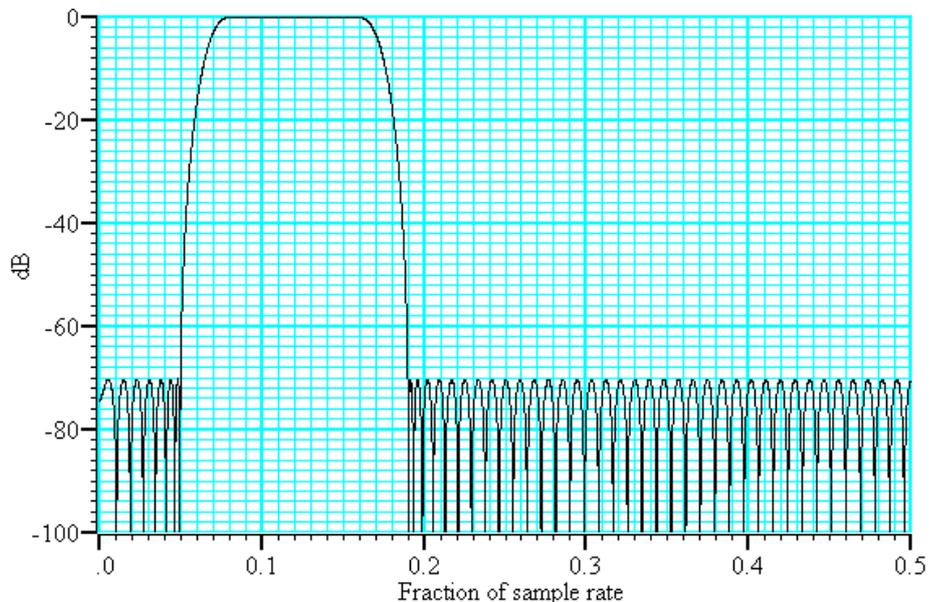
You can define up to 10 bands. However, it is unusual to need more than three. The most common cases with three bands are called band pass and band stop filters. In a band pass filter, you set a range of frequencies in which you want the signal passed unchanged, and set the frequency region below and above the band to pass zero. In a band stop filter you define a range to pass zero, and set the frequency ranges above and below to pass 1.

You must still allow transition bands between the defined bands, exactly as for the low and high pass filters, the only difference is that now you need two transition bands, not one. Also, if you want a non-zero response at the Nyquist frequency, you must have an odd number of coefficients.

For our example we will take the case of a signal sampled at 250 Hz. We want a filter that passes from 20 to 40 Hz (0.08 to 0.16) with transition regions of 7.5 Hz (0.03). If we say it is 10 times more important to have no signal in the stop band than ripple in the pass band, and we want 70 dB cut in the stop band we will get 50 dB ripple in the pass band (because a factor of 10 is 20 dB). To use the formula for the number of coefficients we need the mean attenuation/ripple in dB and the width of the transition region. The two stop bands have an attenuation of 70 dB and the pass band has a ripple of 50 dB, so the mean value is $(70+50+70)/3$ or 63.33 dB. We have two transition regions (both the same width). In the general case of transition regions of different sizes, use the smallest transition region in the formula. Plugging these values into the formula predicts 113 coefficients, however only 111 are needed to achieve 70 dB.

```
var prm[5][3];           ' 3 bands for band pass
var coef[111];          ' 111 coefficients needed
'   band start         band end         function         weight
prm[0][0]:=0.00; prm[1][0]:=0.05; prm[2][0]:=0.0; prm[3][0]:=10.0;
prm[0][1]:=0.08; prm[1][1]:=0.16; prm[2][1]:=1.0; prm[3][1]:= 1.0;
prm[0][2]:=0.19; prm[1][2]:=0.50; prm[2][2]:=0.0; prm[3][2]:=10.0;
FIRMake(1, prm[0][], coef[0]);
```

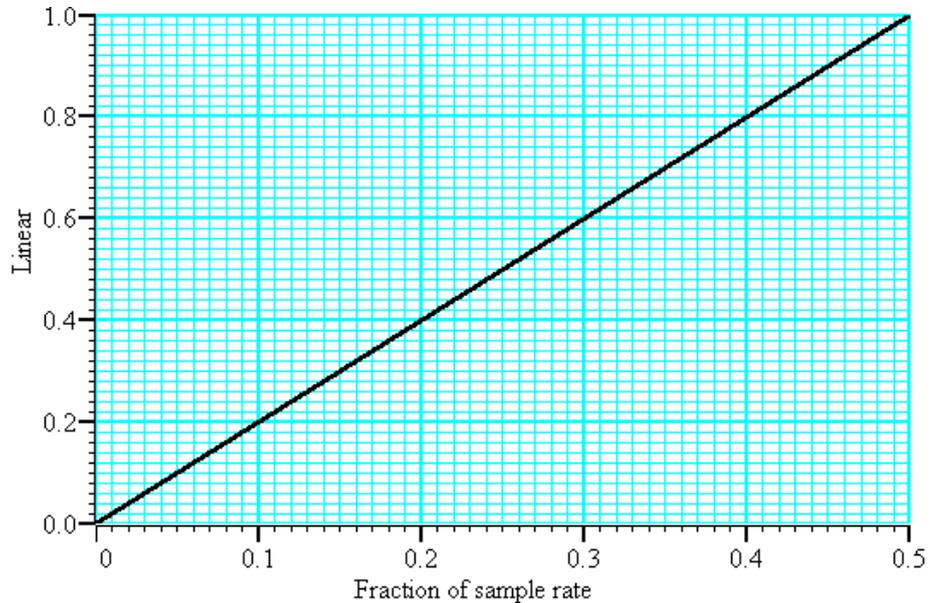
Band pass filter with 111 coefficients



Differentiators

A differentiator filter has a gain that increases linearly with frequency over the frequency band for which it is defined. There is also a phase change of 90° ($\pi/2$) between the input and the output.

*Ideal differentiator
with slope of 2.0*



You define the differentiator by the number of coefficients, the frequency range of the band to differentiate and the slope. The example above has a slope of 2. Within each band (normally only 1 band is set) the program optimises the filter so that the amplitude of the ripple (error) is proportional to the response amplitude. A differentiator is normally defined to operate over a frequency band from zero up to some frequency f . If f is 0.5, or close to it, you must use an even number of coefficients, or the result is very poor. You can estimate the number of coefficients required with the following function:

```
' dB    the proportional ripple expressed in dB
' f     the highest frequency set in the band
' even% Non-zero if you want an even number of coefficients
func NCoefDiff(dB, f, even%)
if (f<0) or (f>0.5) then return 0 endif;
f := 0.5-f;
var n%;
if (even%) then
    n% := (dB+43.837*f-35.547)/(0.22495+29.312*f);
    n% := (n%+1) band -2;    'next even number
else
    if f=0.0 then return 0 endif;
    n% := dB/(29.33*f);
    n% := n% bor 1;        'next odd number
endif;
return n%;
end
```

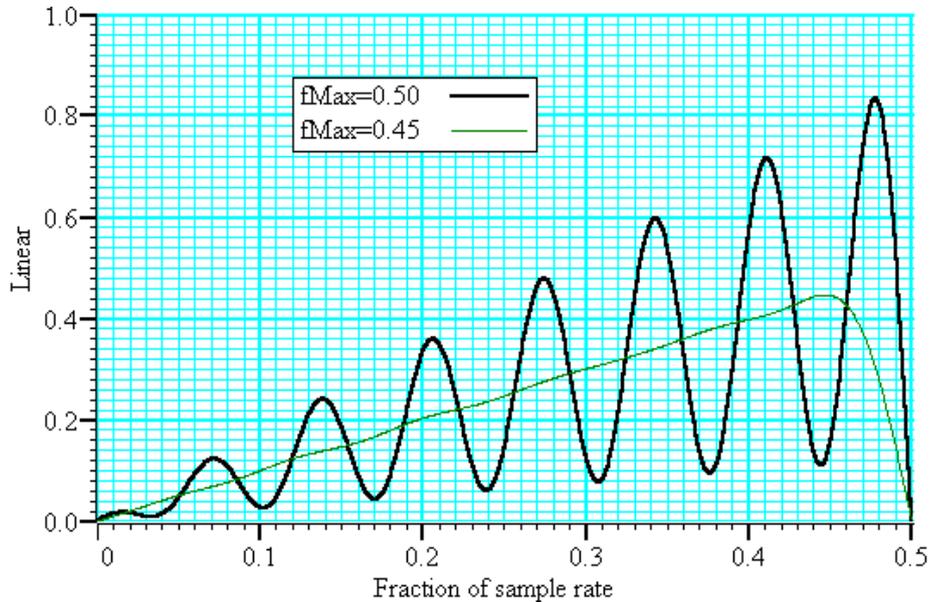
For an even number of coefficients this is unreliable when f is close to 0.5. For an odd number, no value of n works if f is close to 0.5.

These equations were obtained by curve fitting and should only be used as a guide. To make a differentiator that uses a small number of coefficients, use an even number of coefficients and don't try to span the entire frequency range. If you cannot tolerate the half point shift produced by using an even number of coefficients and must use an odd number, you must set a band that stops short of the 0.5 point. Remember, that by not specifying the remainder of the band you have no control over the effect of the filter in

the unspecified region. However, for an odd number of points, the gain at the 0.5 point will be 0 whatever you specify for the frequency band.

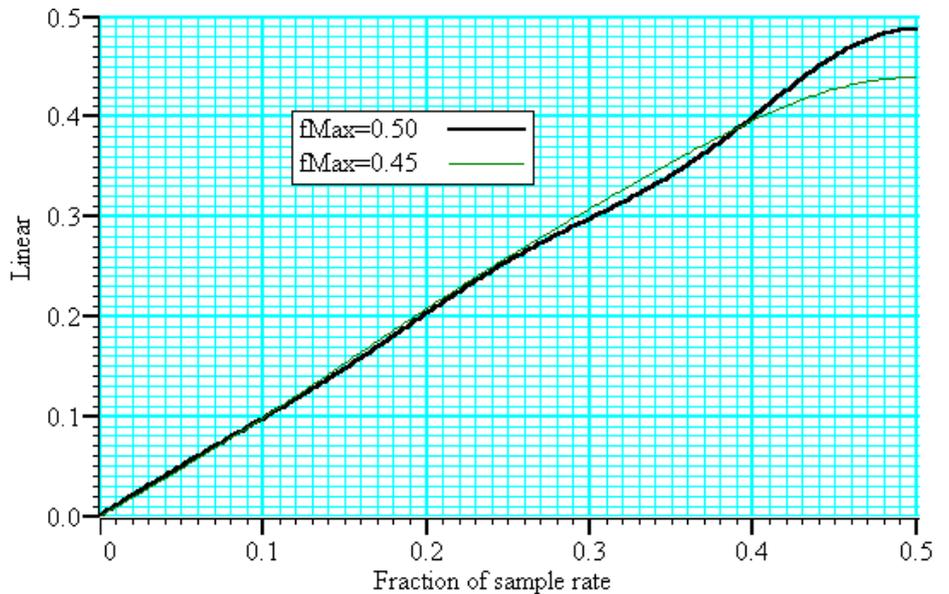
The graph below shows the effect of setting an odd number of coefficients when generating a differentiator that spans the full frequency range. The second curve shows the improvement when the maximum frequency is reduced to 0.45.

Differentiators with 31 coefficients



If you must span the full range, use an even number of coefficients. The graph below shows the improvement you get with an even number of coefficients. The ripple for the 0.45 case is about the same with 10 coefficients as for 31.

Differentiators with 10 coefficients



```
var prm[4][1];      ' 1 bands for differentiator
var coef[10];      ' 10 coefficients needed
' band start      band end      slope      weight
prm[0][0]:=0.00;  prm[1][0]:=0.45;  prm[2][0]:=1.0;  prm[3][0]:=1.0;
FIRMake(2, prm[1][], coef[1]);
```

Hilbert transformer

A Hilbert transformer phase shifts a band of frequencies from F_{low} to F_{high} by $-\pi/2$. The target magnitude response in the band is to leave the magnitude unchanged. F_{low} must be greater than 0 and for the minimum magnitude overshoot in the undefined regions, F_{high} should be $0.5 - F_{\text{low}}$. The magnitude response at 0 is 0, and if an odd number of coefficients is set, then the response at 0.5 is also 0. This means that if you want F_{high} to be 0.5 (or near to it), you must use an even number of coefficients.

There is a special case of the transformer where there is an odd number of coefficients and $F_{\text{high}} = 0.5 - F_{\text{low}}$. In this case, every other coefficient is 0. This is no help to Spike2 and the MSF and MSF programs, but users who write their own software can use this fact to minimise the number of operations required to make a filter.

It is extremely unlikely that a Hilbert transformer will be of any practical use in the context of Spike2, so we do not consider them further. You can find more information about this type of filter in *Theory and Application of Digital Signal Processing* by Rabiner and Gold.

Programmable signal conditioners

Overview

The Spike2 signal conditioner control panel supports the CED 1902 Mk III and IV and the Axon Instruments CyberAmp signal conditioners and also the Power1401 gain option. You can open the conditioner control panel from either the sampling configuration channel parameters dialog (when the channel type is waveform or WaveMark) or from the **Sample** menu. The 1902 and CyberAmp signal conditioners are controlled through a serial port which is set in the **Edit menu Preferences...** option.

What a signal conditioner does

A signal conditioner takes an input signal and amplifies, shifts and filters it so that the data acquisition unit can sample it effectively. Many input signals from experimental equipment are too small, or are masked by high and or low frequency noise, or are not voltages and cannot be connected directly to the 1401.

Signal conditioners may also have specialist functions, for example converting transducer inputs into a useful signal, or providing mains notch filters. The CED 1902 has options for isolated inputs and specialised front ends include ECG with lead selection, magnetic stimulation artefact clamps and EMG rectification and filtering. The only option for the Power1401 is Gain. You should consult the documentation supplied with your signal conditioner to determine the full range of capabilities.

Serial ports

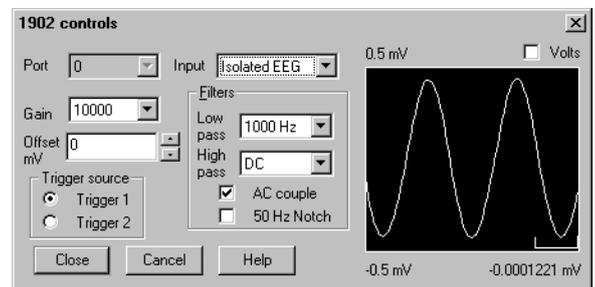
The basic communication parameters are set in the **Edit menu Preferences** dialog. Most supported programmable

signal conditioners are controlled through communication (serial) ports. No port is used if the conditioner support is not loaded or **None** is selected in the preferences or for the Power1401 gain option. Check the *Dump errors...* box to write diagnostic messages to CEDCOND.LOG in the current folder. Do not check the box unless you are having problems with a conditioner as it slows Spike2 down.



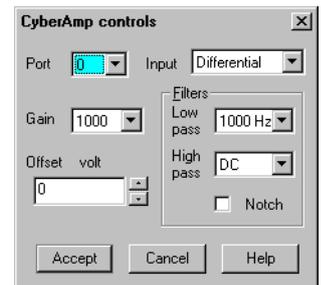
Control panel

The control panel is in two halves. The left-hand half holds the controls that change the conditioner settings, the right-hand half displays data from the conditioner. The right-hand half is omitted if Spike2 is sampling data, or if the 1401 is not available for any other reason.



If the right-hand half is present, the Volts check box causes the data to be displayed in Volts at the conditioner input in place of user units as defined by the Channel parameters dialog. The number at the bottom right is the mean level of the signal in the area marked above the number.

Signal conditioners differ in their capabilities. Not all the controls listed below may appear for all conditioners. This example is the CyberAmp control panel (with the right-hand half omitted). The controls are:



Port This is the physical 1401 port that the conditioner is attached to. If you open the conditioner dialog from the channel parameters dialog you cannot change the port.

Input If your signal conditioner has a choice of input options, you can select the input to use with this field. The choice of input may also affect the ranges of the other options.

Gain This field sets the gain for the signal selected by the **Input** field. Spike2 tracks changes of gain (and offset) and changes the channel gain and offset in the sampling configuration to preserve the y axis scale. You should adjust the gain so that the maximum input signal does not exceed the limits of the data displayed on the right of the control panel. When Spike2 is sampling, the gain and offset are fixed once you have written data to disk.

This is the only editable field for the Power1401 ADC Gain option.

Offset Some signals are biased away from zero and must be offset back to zero before they can be amplified. If you are not interested in the mean level of your signal, only in the fluctuations, you may find it much simpler to AC couple (1902) or high-pass filter (CyberAmp) the signal and leave the offset at zero. If Spike2 is sampling you cannot change the offset once data has been written to disk.

Low-pass filter A low-pass filter reduces high-frequency content in your signal. Filters are usually specified in terms of a corner frequency, which is the frequency at which they attenuate the power in the signal by a factor of two and a slope, which is how much they increase the attenuation for each doubling of frequency. Sampling theory tells us that you must sample a signal at a rate that is at least twice the highest frequency component present in the data. If you do not, the result may appear to contain signals at unexpected frequencies due to an effect called aliasing. As the highest frequency present will be above the corner frequency you should sample a channel at several times the filter corner frequency (probably between 3 and 10 times depending on the signal and the application).

You can choose a range of filter corner frequencies, or you can choose to have the data unfiltered (for use when the signal is already filtered due to the source).

High-pass filter A high-pass filter reduces low-frequency components of the input signal. The high-pass filters area is specified in the same way as low-pass filters in terms of a corner frequency and a slope, except that the slope is the attenuation increase for each halving of frequency. If you set a high-pass filter, a change in the mean level of the signal will cause a temporary change in the output, but the output will return to zero again after a time that depends on the corner frequency of the filter. The lower the corner frequency, the longer it takes for mean level change to decay to zero.

Notch filter A notch filter is designed to remove a single frequency, usually set to the local mains power supply (50 Hz or 60 Hz, depending on country).

The remaining options are for the 1902 only:

AC couple This is present for the 1902 only, and can be thought of as a high-pass filter with a corner frequency of 0.16 Hz. However, it differs from the high-pass filters as it is applied to the signal at the input; the high-pass filters in the 1902 are applied at the output.

Trigger source The 1902 provides two conditioned trigger inputs, and one output. This control selects which of the inputs is connected to the output.

Filter type A 1902 mk IV with digital filters has extra fields that allow you to set the low-pass and high-pass filter types. You can choose Butterworth or Bessel characteristics and 2 pole or 3 pole filters.

Setting the channel gain and offset

If you change the gain or offset in the control panel, Spike2 will adjust the channel gain and offset in the sampling configuration to compensate so as to keep the y axis showing the same units. This means that if you change the gain, the signals will still appear in the same units in the file. However, the first time you calibrate the channel you must tell the system how to scale the signal into y axis units.

For example, to set up the y axis scales in microvolts you do the following:

1. Open the Sampling Configuration dialog.
2. Select the waveform or WaveMark channel and open the channel parameters dialog.
3. Set the Units field of the Channel parameters to uV.
4. Set the Input in Volts x field to 1000000.
5. Press the Conditioner button to open the conditioner control panel.
6. Adjust the gain to give a reasonable signal.
7. Close the signal conditioner control panel.

You only need do steps 3 and 4 once. Any subsequent change to the conditioner gain will adjust the channel gain to leave the units in microvolts.

For the more general case where you have a transducer that measures some physical quantity (Newtons, for example) and it has an output of 152.5 Newtons per mV. If you wanted the Y axis scaled in Newtons, you would replace steps 3 and 4 above with:

3. Set the Units field of the Channel parameters to N.
4. Set the Input in Volts x field to 0.1525.

To work this out you must express the transducer calibration in terms of Units per Volt (in this case Newtons per Volt).

If you have set an offset in the conditioner, and you want to preserve the mean signal level, you should null it out by changing the offset in the Channel parameters dialog.

Conditioner connections

Spike2 can control multiple signal conditioners, but they must all be of the same type (selected when you install Spike2). Each conditioner has a unit number starting at 0 and each unit conditions the same number of data channels (1 for 1902, 8 or 2 for the CyberAmp, 8 for the Digitmer 360). For both the 1902 and the CyberAmp, if the conditioner unit number is u and there are n channels per conditioner, it conditions 1401 ADC channels $u*n$ to $u*(n+1) - 1$. It is up to you to make the correct connections.

Basic communication and connection information is stored in the registry in the key `HKEY_CURRENT_USER/Software/CED/CEDCond`. For each conditioner type that needs to store information there is an additional sub-key, currently 1902, CyberAmp, D360, Power1401. This sub-key holds integer variables:

- Port** The communication port number in the range 1 to 19. The default value is 1.
- Dump** Optional. Set to 1 to write the file `CECOND.LOG` holding diagnostic messages to the folder that Spike2 ran from.
- First** Optional, default 0. This sets the first 1902 or CyberAmp unit number to search for. Use `First` to save time waiting for missing units to time out.
- Last** Optional, default 0. Searches stop when there is no response from a unit number greater than `Last`. Use `Last` to skip over missing units.

The `Port` and `Dump` values can be set in the Edit menu Preferences option. Before revision 7.09 this information was saved in the file `CECOND.INI` in the system folder. If the registry does not contain the information, we read the `.INI` file. The file format is:

```
[General]
Port=COM1
Dump=1
First=1
Last=3
```

CED 1902 The CED 1902 unit numbers are set by an internal switch pack; multiple units usually have the channel number as a label on the front panel. Each 1902 output should be connected to the 1401 ADC input with the same number as the 1902 unit.

Usually the 1902 unit numbers start at 0 and are contiguous, in which case you do not need to set `First` and `Last`. As an example of a more complicated situation, let us suppose you have unit numbers 4, 5, 6, 7, 12, 13, 14 and 15. In this case, you should set `First` to 4 and `Last` to 15 for ADC channels 4-7 and 12-15. The Try1902 test program has options to set `First` and `Last`.

CyberAmp The CyberAmp has a device address (unit number) set by a rotary switch on the rear panel; multiple units must have different addresses. There are two types of CyberAmp: 8-channel and 2-channel. If you have multiple units, they must all have the same number of channels, or all the 8-channel units must have lower device addresses than all the 2-channel units.

Normally you will have one 8-channel CyberAmp, and you will give it address 0 to support ADC channels 0 to 7. In this case you do not need to set `First` and `Last`. If you want to connect it to channels 8-15, you would set both `First` and the device address to 1 and then you could connect it to channels 8-15.

Here is a more complex example with two CyberAmp 8-channel units and one 2-channel. The table shows some possible configurations (assuming 32 ADC inputs):

CyberAmp configurations

CEDCOND.INI		8-channel unit 1		8-channel unit 2		2-channel unit 3	
First	Last	Switch	ADC	Switch	ADC	Switch	ADC
0	2	0	0-7	1	8-15	2	16-17
1	3	1	8-15	2	16-23	3	24-25
2	3	2	16-23	3	24-31	-	-

Digitimer D360 Install the Digitimer-supplied D360 support first and verify that it is working by running the Digitimer D360 Client application. The Digitimer D360 Client software must be at least version 4.7.0.0; upgrade to the latest version if it is older. The D360 connects to the computer with a serial line (COM port), however, our code has no control over the port; this is configured by the Digitimer code. The only registry option is `Dump`.

Multiple D360 amplifiers can be used; each has 8 channels numbered 1 to 8. Channels 1 to 8 of the first D360 are assumed to be connected to ADC inputs number 0 to 7, the second D360 drives ADC ports 8 to 15 and so on. The amplifier numbers are determined by the order in which they are daisy-chained together, they should be connected in serial-number order with the lowest-numbered amplifier connected directly to the PC, the next lowest-numbered amplifier connected to the first amplifier and so on.

Power1401 with programmable gain Spike2 will automatically detect the channels. The only registry option is `Dump`.

1401-18 Programmable Discriminator

Spike2 can control a CED 1401-18 programmable discriminator card fitted in a 1401*plus*. The **Sample** menu **Discriminator Configuration...** command opens the discriminator setup tabbed dialog. The dialog has sections that control the signal routing through the card, the mode of discrimination, the trigger levels and the waveform channel used as a signal monitor.

What the 1401-18 does

The 1401-18 event discriminator card detects waveform events using two trigger levels and an optional time window. The output from the detector is a digital signal suitable for input to the Spike2 event channels as event and level event data.

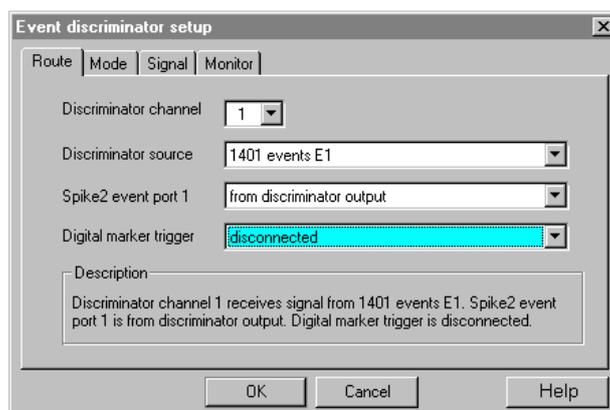
The card contains eight discriminator channels numbered 0 to 7. Channels 0 to 5 can take their input from either a front panel event input BNC or a digital input. Channels 6 and 7 take their input only from the digital input port. The input voltage range is ± 5 Volts. The output of the channels is normally routed to the Spike2 event channels, however channel 1 can be routed to the digital marker trigger and channel 3 to the start sampling trigger.

Spike2 saves the discriminator setup as part of the sampling configuration. The dialog is used before sampling starts to prepare the 1401-18; it cannot be used during sampling. However, you can use the script language discriminator support during sampling.

Signal route

The **Route** tab in the Event discriminator setup dialog sets the source and destination of the discriminator, Spike2 event port, digital marker trigger and sampling trigger data. All eight channels are independent of each other.

The text **Description** field gives a synopsis of the settings for the channel. The fields are:



Discriminator channel

This selects the discriminator channel to adjust. It also sets the channel for the **Mode** and **Signal** tabs. You can also change channel in the **Signal** tab.

Discriminator source

This selects the input to take waveform data from. All channels can take data from the digital input connector, and channels 0 to 4 can use one of the front panel event inputs. Channel 5 can take its input from the ADC Ext front panel input. You can also set the source as **Not used** in which case the signals for this channel behave as if the discriminator were not installed. When a front panel input is used as a source, that input cannot be used as the source of any other signal.

Spike2 event port n

This field sets where Spike2 takes the data used for Event and Level data from for port n. This can come from the digital input, or from the discriminator output. You can also disconnect it (in which case you will not be able to collect data from this port).

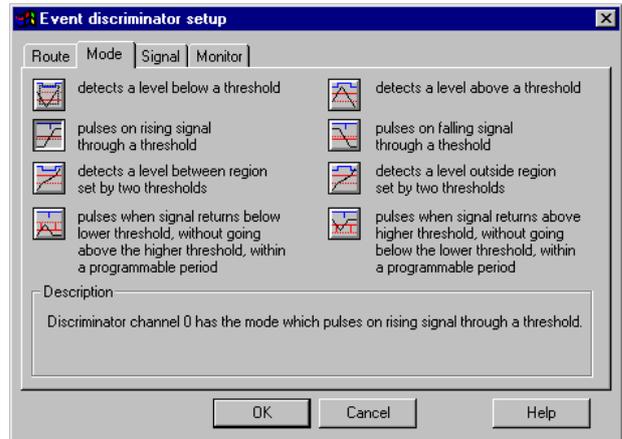
Digital marker trigger Sampling trigger

The **Digital marker trigger** field is present for channel 1 only, the **Sampling trigger** field is present for channel 3 only. You can choose to take these signals either from the front panel E1 (**Digital marker trigger**) or E3 (**Sampling trigger**) inputs, from the discriminator output, or to leave the trigger disconnected.

Discrimination mode

The Mode tab in the Event discriminator setup dialog sets the discrimination mode for the current channel and gives you a description of each mode. You can change the channel in the Route and Signal tabs. You can also set the mode from the Signal tab.

Set the mode by clicking one of the buttons. There are eight buttons, but only four modes; the right-hand column modes are the same as the left-hand column but with the active signal in the opposite direction.



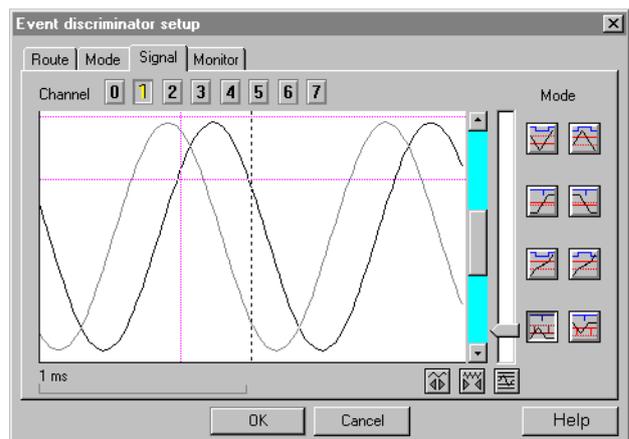
In the pictures on the buttons, the main threshold is shown as a solid line. The dotted line is a second threshold. The four modes are:

1. The first mode detects the signal above or below a band defined by the two levels. When the signal lies between the levels, the last level crossed determines the state of the output. The output is suitable for use in a Level data channel.
2. The second mode produces a 1 μ s pulse when the signal crosses the main threshold level. The data must cross the second threshold before another pulse can occur. The output is suitable for an Event data channel.
3. The third mode detects when the signal lies between the two thresholds (or conversely, lies outside the levels). The output is suitable for use in a Level data channel.
4. The last mode produces a 1 μ s pulse when the signal crosses the main threshold and falls below it again without crossing the second threshold within a set time period. The output is suitable for an Event data channel.

View data

The Signal tab shows the raw and discriminated waveform on a selected channel so you can set the threshold levels. The buttons along the top of the window let you swap channel quickly and the buttons on the right swap modes (see the Mode tab for descriptions of each mode).

The vertical scroll bar moves the data vertically in the window. The slide bar to the right of the scroll bar changes the vertical magnification of the displayed wave. If your signal looks very small on the screen even with the gain slider at the top, your signal needs more amplification before it can be usefully discriminated by the 1401-18. At maximum gain, the threshold levels are accurate within a few pixels in the vertical direction.



The two buttons below the vertical scroll bar change the time width of the displayed data. The button under the slide bar optimises the display and threshold levels based on the current trigger mode and recent data.

To set up a channel, follow these steps:

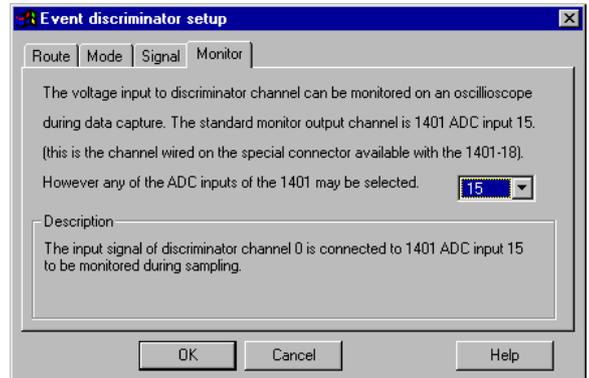
1. Select the channel with the buttons above the displayed data.
2. Select the discrimination mode with the buttons on the right.
3. Use the width controls and the optimise button to display data and cursors.
4. Adjust the trigger levels (and the time window if enabled) by dragging the trigger levels (and time marker) with the mouse.
5. Repeat for all channels and click OK when you have finished.

The data area shows the last event that caused a trigger and the current input (if there is no triggering event). During setup, the computer may not be fast enough to show you every event that causes a trigger. However, during sampling, all such events are captured.

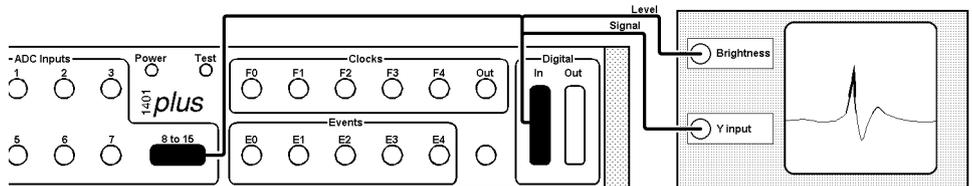
You may sometimes see trigger events that do not appear to cross a threshold. This can happen when the data contains spikes of very short duration. The displayed data is sampled, and it can happen that very short spikes fall between samples, so are invisible on the screen. This is usually an indication that the input data contains high-frequency noise and should be passed through a filter.

Monitor channel

Spike2 uses a waveform input channel to monitor the data entering the discriminator. The cables supplied with the 1401-18 card connect the monitor output to ADC channel 15. It is unusual to change the monitor channel. The supplied cables also have labelled BNC connectors for an oscilloscope. The oscilloscope outputs are active when Spike2 captures data (when the Signal option cannot be used) and can be used to monitor the discriminator action on the current channel.



Cable connections with the 1401-18



It is important to fit this cable (or make the connections with your own wiring) so that the Signal option can display the current discriminator channel. If you connect your input signals to the 1401 through the digital inputs, you must open the 25-way connector on the cable and wire your signals into it.

Digital input connections

The functions of the digital input connector are changed with the 1401-18 card present. The pins of the digital inputs that previously were the Spike2 event inputs (see the *Sampling data* chapter) now become discriminator inputs when selected in the Route tab. The discriminator outputs are buffered and connected to the pins that were previously the digital marker inputs (TTL output pin in the table below). You can still use digital markers as the digital marker data bit inputs are duplicated on the same pin numbers on the digital output port. However, the analogue ground uses pins 23 and 24, so the optional strobe and h/s signals are not available if you have the discriminator fitted.

Digital input pins for discriminator channels

Discriminator channel	7	6	5	4	3	2	1	0
Data input pin	1	14	2	15	3	16	4	17
TTL output pin	5	18	6	19	7	20	8	21

Other digital input signal pin numbers

Analogue ground	22, 23 and 24			Signal Monitor output	11
Level Monitor output	10			+5 Volts	25
Digital ground	13				

The Signal Monitor output is the raw waveform data for the current discriminator channel. The Level Monitor output signal is a 3-level signal that can be used for 'scope brightness controls, indicating where the input lies with respect to the two threshold levels.

Electrical information

All 1401-18 waveform inputs have a working range of ± 15 Volts. There is a 20 Volt over-voltage protection on all inputs, even on power-off. When the event discrimination is by-passed, the card has a maximum impedance of 50 Ohms. When a channel is set for conditioning, the impedance is 100 kOhms, with 100 nA bias current and up to 2 mV offset.

The two thresholds have a ± 5 Volt programmable range, with a 12-bit resolution. Linearity is to ± 2 bits, the offset accuracy is ± 5 bits. There are no adjustments.

There is a fixed 20 mV hysteresis in all modes built-in to the hardware. In modes that use the second threshold level to program hysteresis the software allows for the extra 20 mV. This hysteresis is present to stop noise on the input causing spurious, high frequency triggering.

Input pulse widths should be at least 1.2 μ s; pulses narrower than this may not be detected. Pulse outputs are approximately 1 μ s in length.

The time window is programmable from 20 μ s to 0.65535 seconds in steps of 10 μ s.

The two monitor output signals can drive loads as low as 600 Ohms. The Level Monitor output signal has the following voltage levels:

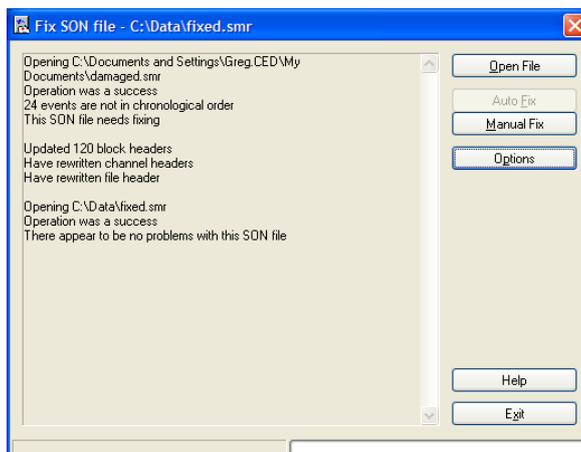
2 to 5 Volts	Signal above both thresholds
1 to -1 Volts	Signal between thresholds
-2 to -5 Volts	Signal below both thresholds

Utility programs

The SonFix data recovery utility



The SonFix utility program recovers damaged Spike2 data files. A Spike2 data file has a header followed by a channel description list, followed by data blocks. Each data block has links to the previous and next data block of the same channel. If a file becomes damaged, these links can be broken, rendering the file useless. The SonFix program scans the file, rebuilds the links for each channel and corrects the file information in the file header. There are two common types of damage that occur to Spike2 data files:



1. The file is damaged as it is sampled. This can happen if your system loses power. The next time you run Spike2 it will tell you where the data can be found and recommend that you run SonFix to repair the file.
2. The file is damaged due to some other disk accident (such as a disk crash or when archiving and restoring files to tape, CD, DVD or some other removable media). As long as the damage is not at the start of the file, the file can usually be fixed and undamaged data recovered.

This version of SonFix relies on the header and channel list being intact. However, if the program does not recognise the file as a Spike2 data file, don't despair. It is possible to patch the file header with a header from a similar file, and then recover the data. Contact the CED Help desk if you are in this situation.

Using SonFix

If you have a disaster, check that your disk drive is not damaged. Writing more data to a damaged drive is likely to cause further problems and may render your data unrecoverable. If the hard disk is damaged you must repair it with a suitable disk utility program or find someone who can do this for you before running SonFix.

Next, find the data file. If you have to use a disk repair utility, you may find it useful to know that the first two bytes of a data file hold the file revision followed by the ASCII characters "(C) CED 87".

In most cases there will be no disk damage and the file will be present with a non-zero size. **If your file is very valuable**, make certain you do not fix it in place by saving the fixed file with the same name as the damaged file. Run SonFix by double clicking its icon (it is in the same folder as the Spike2 program). Click the **Open File** button and select the file to repair and SonFix will tell you if it can repair the file. Click the **Auto Fix** button to carry out the repair.

The **Options** button opens a dialog that controls the tests carried out on a data file to detect damage and sets how event and marker channels are checked and repaired. If the data file is heavily damaged, it may be necessary to use the **Manual Fix** option. Contact CED if you have problems fixing a file or using the manual fix tools.

What cannot be recovered

Data that was not written to the physical disk cannot be recovered. To keep the system responsive and to make high speed sampling efficient, Spike2 buffers data in memory as much as it can. The operating system also buffers data in memory before writing it to disk. However, with data in memory, a power loss or system crash will lose data.

If you use the **Flush data to disk** option in the sampling configuration, you can protect yourself against disaster at the cost of some loss of performance. If you do not use the flush data option you increase the risk of data loss if your system goes down.

Batch Processing

There are two ways to test or fix more than one file at a time:

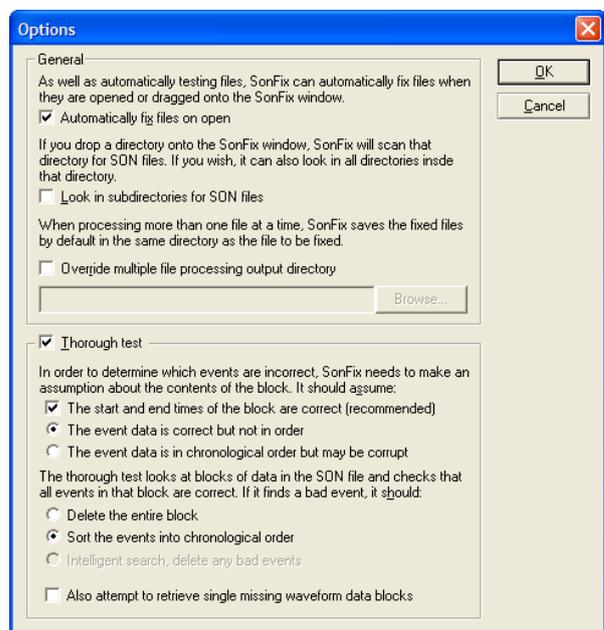
1. Use the **Open File** button on the main window to show a standard File Open dialog box in which you can select as many SON files as you wish to test or fix.
2. Drag files from Windows Explorer or My Computer and drop them on the SonFix window is the other way of testing or fixing several files at the same time. You can also drag a directory onto the SonFix window, and it will test or fix all the SON files (files with a .smr extension) in that directory, and optionally all the SON files in any sub-directories (this setting can be changed from the Options window).

If you want files to be fixed automatically, you can specify this in the **Options** window. Fixed files have **Fixed** added on to the start of the file name. The fixed files will, by default, be placed in the same directory as the damaged files, but this can be changed from the **Options** window.

After SonFix finishes testing and optionally fixing your data files, it displays a list of the files that needed fixing, and a summary of the number of files tested and fixed.

SonFix Options

To customise SonFix, select **Options** from the main SonFix window. As well as the options described above, there is also a section in the options dialog box entitled **Thorough Test**, which enables a complete test of the times of event and marker data. If selected, the thorough test is performed on files after the main test, and checks all events in the data file to make sure that they are in the correct order. To help with the error detection, you can select various assumptions that SonFix should make when it attempts to recover data from a damaged file:



Block start and end times are correct

You should normally select this option. It tells SonFix to assume that only the actual event data is corrupt, not the data block headers. If, when you use this option, you find that a great deal of event or marker data is deleted, it would be worthwhile trying to fix the file without this option to see if that helps.

Event data is correct but not in order
Event data is in order but may be corrupted

You can choose between one of these two possibilities, which determine both how the event data is checked and also how it may be fixed. There is no correct choice here; you may simply have to try both options to see which gives the better fix. To start with, selecting data in order but corrupted is probably a better guess.

If SonFix detects event or marker data corruption, it can fix problems in these ways:

- by deleting the block containing the corrupt events (a good idea for heavily damaged files where you care more about the waveform data than event or marker data), or if the other repair option fail.
- by sorting the events into numerical order or, in the case where two events have the same time, adjusting the time of the second event to be just later (for example for a file containing RealMark, WaveMark or TextMark channels where you care more

about the data they contain than the times they are stored at). This repair mechanism is only available if you are assuming that event data is out-of-order but not corrupt.

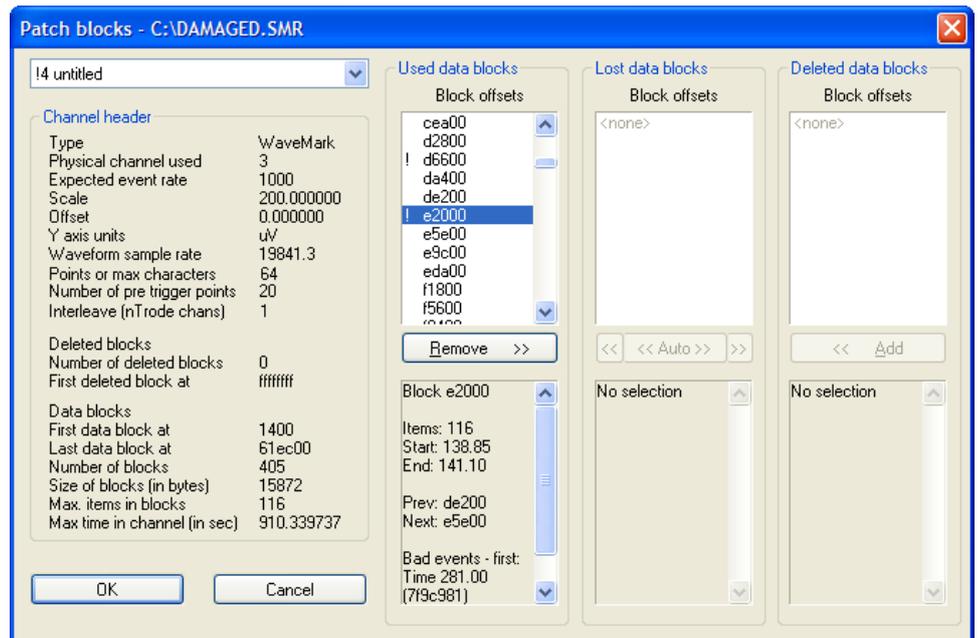
- by doing an intelligent search for bad events and deleting the minimum number of events to leave the data in order. This is usually the best option. It is only available if you are assuming that event data is in order but corrupted.

Also attempt to retrieve missing single waveform data blocks

If a waveform data block has a corrupted header, it may not appear in any block list. If the used list has a gap where the forward and backward pointers both agree on the position of a missing block, and no other channel claims it, SonFix can recreate the block header. This is possible because waveform data is sampled at a known rate; this cannot be done for event or marker data. This option takes advantage of this to retrieve single missing blocks of waveform data. Retrieved data blocks may contain bad waveform data, but recovery removes gaps in the channel, which could be useful.

Manual Fix

The Manual Fix button leads to the Patch blocks dialog in which you can inspect the lists of used, deleted and lost data blocks for each channel and other channel details. You can also move blocks between these three lists for each channel. Unless you are certain that you know what you are doing it is best to only use Auto Fix; the manual fix tools are intended for use primarily by CED personnel or under CED direction.



The Patch blocks dialog shows information for one channel at a time; the channel selector is at the top left of the dialog. Channels with detected problems have an exclamation mark before the channel number. You should be able to ignore channels without such an indicator.

The channel header information is displayed on the left side of the dialog. The information shown includes general channel information such as the channel type and sampling rate, information about any deleted blocks (these are disk blocks that held data when the channel was deleted at some point in the past, and have not yet been reused) and details of the channel data blocks themselves.

The detailed data block information is shown at the right of the dialog. All data blocks are multiples of 512 bytes in size. When SonFix scans the data file, it searches for data that looks like the start of a block at each 512 byte boundary in the file. The block headers hold the channel number to which the block belongs. When it finds a block, it

can use the block size information in the file header to predict where the next block could start. The scan of the file allows Sonfix to build up block lists for each channel:

- Used data blocks* Used data blocks are blocks of data that are correctly linked into the channel and that do not show any problems. For a channel that is OK, all of the data blocks will be shown in the Used list (with possibly some in the deleted list).
- Lost data blocks* Lost data blocks are blocks that apparently make up part of the channel data (because they have this channel number in the block header) but which are not connected properly to the rest of the channel data or the header information does not make sense. Fixing a channel usually comes down to moving these lost blocks into either the **Used** list or into the **Deleted** list.
- Deleted data blocks* Deleted blocks are blocks forming part of the deleted channel data. When you delete a channel in a data file, the chain of blocks belonging to the channel are placed in the channels deleted block list. If you reuse the channel with the same block size, the deleted blocks are reused. If you re-use the channel with a different block size, the deleted blocks are abandoned. The presence of deleted blocks does not indicate a problem.

Each set of information shows a list of data blocks above, and information about a single (selected) data block below. The list of data blocks shows each data block's address within the disk file as a hexadecimal (base sixteen) number. Blocks that are deemed invalid (either in relationship to the adjacent blocks in the list or as a result of internal consistency checks) are shown with a preceding exclamation mark, blocks that make up part of a coherent chain of blocks descending from an initial block are shown indented to the right with the initial block not being indented. The block information shown in the area below shows the number of items, the time range covered by the block, the previous and next blocks in the chain and details of any errors detected.

Use the buttons at the bottom of the block lists to move blocks between lists. The **Remove** button below the **Used** list moves the selected block into the deleted list, while the **Add** button in the deleted list moves a block into the Used list. However these options are rarely required compared to the buttons below the **LOST** list. The button marked '<<' moves a block, or set of blocks, into the **Used** list. Normally it will only move the selected block but, if this block is the start of a coherent set of blocks all shown indented to the right, then all of the blocks within this coherent set are also moved. The button marked '>>' does the same thing, but it moves blocks into the **Deleted** list.

The third button is marked '<< Auto >>'. This does the following:

1. Attempts to move all of the lost blocks into the used list, as long as this will not result in a corrupt or invalid set of blocks.
2. If step 1 fails, finds the largest single coherent set of blocks in the lost list that can be moved into the used list without problems and moves it.
3. If no blocks can be moved into the used list without problems, all of the lost blocks are moved into the deleted list.

This corresponds roughly to the process carried out by the **Auto Fix** button in the main SonFix window, though the **Auto Fix** button repeatedly executes the three steps shown above until the lost list is empty.

It is not possible to describe how to use the Patch blocks dialog completely as the steps required will vary according to the file damage. Roughly speaking, you should attempt to first find out which lost blocks are hopelessly corrupted and move them to the deleted list, then move all of the others to the used list. Once all of the rubbish is in the deleted list the **Auto** button should do this for you. It is very strongly recommended that you make a backup copy of the damaged data file before attempting a manual fix so that mistakes can be corrected.

Try1401 test program

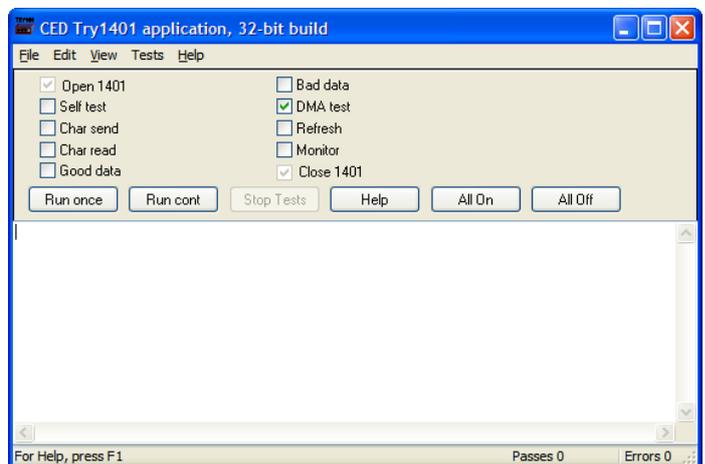
If you are having problems sampling data, and you suspect that the problem lies with the 1401, interface card or device driver, then it is well worth running the Try1401 program, found in the same folder as Spike2 (the program file is called `Try1432.exe` on Windows systems).

The Try1401 program exercises the 1401 in much the same way as Spike2 does, but it also checks every byte of data transferred and reports errors in a way that is useful to CED engineers, particularly when there are data transfer problems. Unfortunately, the sternest test of data transfer we know of is the Spike2 program, so it is possible for Try1401 to pass a system that fails in Spike2 (but this is most unusual).

The Try1401 program can also run the 1401 self-test and decode any errors. These can vary from simple problems with the calibration of the 1401 inputs and outputs (usually not a serious problem), to serious internal errors associated with damaged components. You should select the self-test option if the 1401 flashes the Test LED on power up when no other external equipment is connected to the 1401 inputs.

Running Try1401

Double click the program icon to start the program. You have a choice of test options and of running once or continuously. Errors are written to the text window, and can be copied, cut and pasted. If a single pass of the test doesn't show a problem, run the test continuously and leave it for several hours.



Most of the tests are to check the data path between the host computer and the 1401. If you have a problem with your 1401, the CED engineer who helps you to sort it out will most likely ask you to run this program and email the results. For a confidence check, the only option that is needed is the **DMA test**. The remaining tests help CED engineers to narrow down problems. The individual options are:

Self test This runs the internal 1401 self test and interprets the results. Please remove all connections from the 1401 except the power cord and the data interface cable as other connections can cause self-test failures.

Char send This checks the general data path from the host computer to the 1401. If you have a USB connection, problems indicate some sort of installation failure. For the other interface cards data corruption problems are rare, and indicate either a damaged data cable, or bent pins on the 1401 or the host interface card. A timeout error in this test usually indicates an interrupt problem.

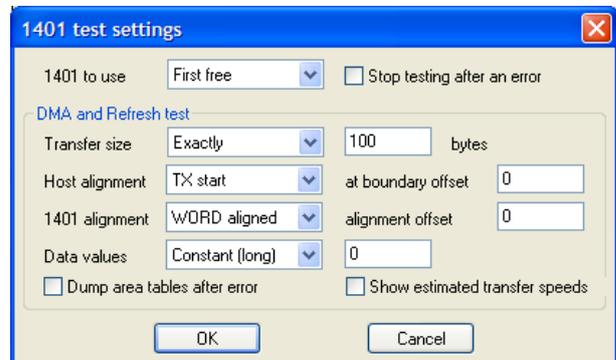
Char read This is similar to *Char send*. It checks the general data path with an emphasis on reading back data.

Good data/Bad data These are further variations on the *Char send* and *Char read* tests. You can usually skip these tests unless a CED engineer asks you to run them.

DMA test This test checks out high-speed block data transfer between the 1401 and your computer. There is no point running this test if the *Char send* or *Char read* tests have failed. This

test transfers data blocks of various sizes and alignments and checks that the data transferred correctly. Unfortunately, Spike2 is the sternest test of data transfers that we know of, but this test will usually detect any transfer problems. There are additional options in the Tests menu DMA settings... command:

You will not normally need to set any of the values in this dialog unless a CED engineer needs additional data to diagnose a data transfer problem. The fields in the dialog are:



Transfer size can be set to: Unconstrained, Exactly, Greater than or Less than. For all sizes except

Unconstrained you must provide a byte count. For all except Exactly mode, Spike2 chooses random sizes within the constraint you have imposed.

The data for transfer is held in memory. This memory lies in a continuous block of virtual address space. However, the physical memory that makes up this address space may be mapped anywhere in computer memory. This usually means that a data transfer is broken up into sub-blocks of contiguous physical memory. Alignment relates to the position of data blocks relative to starts and ends of these sub-blocks. This field only applies to ISA and PCI interface cards using DMA transfers; we have no knowledge of alignment for the USB interface. You can choose from Unconstrained, TX start, TX end, RX start and RX end. TX = transmit to 1401, RX = receive from 1401. The at boundary offset field sets the relative position of the transfer start or end to sub-block boundaries. You can set from -4095 to +4095 (but useful values are usually in the range ± 100).

Data values can be set to be Random, or you can choose to set values based on bytes, words (2 bytes) or longs (4 bytes) and the values can be a constant, an upward ramp or a downward ramp.

Check the Dump area tables after error box to print a lot of extra info about any failing DMA transfer. This can help a CED engineer to diagnose a fault.

Show estimated transfer speeds prints the approximate number of kB per second for each transfer.

File menu There are options in the File menu to update the Power1401 flash memory. These are described in the Power1401 Owner's manual. There is also the 1401 info... command. This prints out information about the 1401 device driver and the 1401, for example:

```
1401 type          = Micro1401
Monitor revision is 20.14
1 megabyte base memory
Micro1401 main card is 2501-01 C-211
Block transfers use DMA, multiple transfer areas
ADC channel sequencer is FIFO at 4 MHz
Supports up to 256 channels, 3uS ADC block
No extra ROM in spare slot
```

Multimedia recording

Multimedia recording



You can use the `s2video` application to record multimedia files automatically whenever you sample a Spike2 data file. You can run up to 4 instances of this application at a time. Each instance will record one video and/or one audio track. The quality of the recording (video resolution and frame rate and audio sample rate) will depend on the hardware and software that you have installed. You can use codecs to compress your multimedia data either on-line during capture or off-line after capture ends.

The multimedia data files are saved alongside the Spike2 data files. You can view these files within Spike2 using the **View** menu **Multimedia Files** command.

File names

Spike2 samples data to a temporary file. When sampling ends and the user chooses to save the file, it is given a permanent name. The multimedia recorder tracks the Spike2 activity and names the data files to match Spike2. For example, if the final Spike2 name is of the form:

```
C:\Spike7\data\data23.smr
```

and if there were two instances of `s2video` running, they would name their files:

```
C:\Spike7\data\data23-1.avi
```

```
C:\Spike7\data\data23-2.avi
```

Availability of files

Multimedia files generated by `s2video` are not available for use until the Spike2 data file has been saved. If they require off-line compression, this is done as a background task so that data capture can continue. In this case, the multimedia files are not available for review until the compression task has been completed.

System requirements

To run Spike2 data capture and multimedia capture simultaneously, you need a suitably powerful computer. The software will take full advantage of multiple processors and hyper-threading. The data capture and replay software is based on Microsoft DirectShow, which is part of DirectX. If you are running Windows XP you should check you have the latest version. You can download this from the Microsoft web site: go to the Downloads section and search for DirectX. You want DirectX 9.0c or later. If you are running Vista or later you should be up to date with the latest versions.

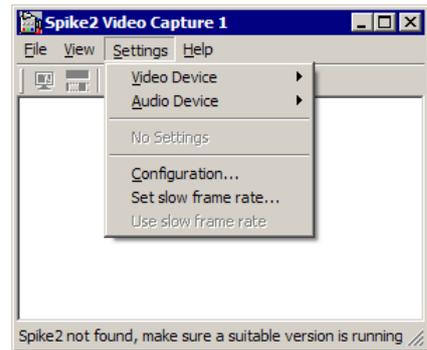
To record video, you will need a web cam or a video camera. There is a very wide variety of these available. You may need to install drivers for your camera; these will be supplied with it. Alternatively, your system may recognise the camera when you plug it in and locate the correct drivers for you. To record multimedia audio you need a sound card in your computer or a video device that also records sound. If you need precise time synchronisation of sound, the best solution is to record it as a Spike2 Waveform channel.

If both your camera/audio device and 1401 link to the computer by USB, and your computer does not have independent USB ports, the two devices will compete for bandwidth. This can degrade the maximum sampling rate of the 1401 and the maximum frame rate of the video.

If your computer has both USB 1 and USB 2 ports, do not plug a video device (or a 1401!) into a USB 1 port. This severely reduces the bandwidth. Even modest video needs a bandwidth of many MB per second and USB 1 will limit you to just over 1 MB per second. For example, my web cam set to 320x240 at 30 frames per second generates 6.4 MB per second. If I plug it into a USB 1 port, it drops more than 80% of the data frames and the result cannot be synchronised.

Getting started

When you run the `s2video` application for the first time, there will be no devices selected. The window title **Spike2 Video Capture 1** means that this is the first instance of the application. If you open another, it will be number 2, and so on. The application stores information in the system registry for each instance. This includes the window position, the video and audio recording devices and any codecs that you select to compress the multimedia data so that it occupies less disk space.



Your first task is to choose the devices to record from. Click the **Settings** menu and then select the **Video Device** item. A new menu will pop-up with a list of the possible video sources plus **No Video Device**, which will have a tick next to it. If a device that you want to use does not appear in the list, make sure that it is switched on and connected. You do not have to select a video device; recording audio only is allowed. Once you select a video device, you will see a preview of the image in the video display area of the window. You can enable and disable the preview with the **View** menu **Preview** command or with the first button in the toolbar. If you select **No Video Device**, the window will reduce in size to hide the video area.

The next task is to select an audio device. Click the **Audio Device** item in the **Settings** menu to display a list of possible audio devices, plus **No Audio Device**. If you selected a video device that generates audio as an integral part of the video stream, the **From Video Device** item is enabled and can be selected. When you select an audio device, an audio level meter appears in the dialog. You do not have to select an audio device; recording video only is allowed.

You must set at least one device to sample data. Once you have set your devices, new items for **Video Device Properties**, **Video Capture** and **Audio Device Properties** are added to the menu.

This software is written to work with a very wide range of hardware devices. We cannot predict what you will see on the screen in the property dialogs for the devices as the device manufacturers define them. We can tell you the types of control to expect from our experiences working with a range of different cameras and sound devices. If you have problems setting up your video or audio device please refer to the documentation that was supplied with them. We will be able to offer general guidance only.

Video Device Properties...

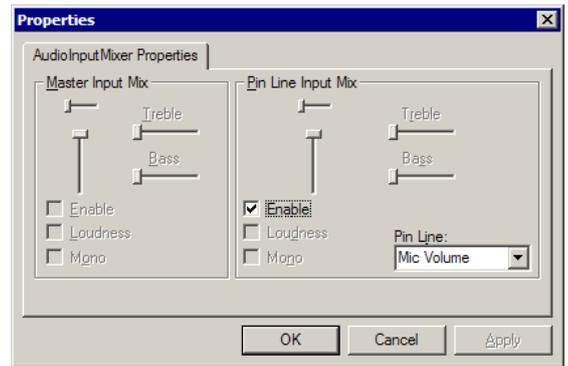
This menu item appears in the **Settings** menu when you have selected a video input. What you see in this menu depends on the selected camera. It will probably include controls for image brightness, saturation, colour balance and saturation. It may also have controls for the frame rate (this will set the maximum frame rate that you can record at), monochrome or colour output and various settings to compensate for background illumination and to cope with fluorescent lighting.

If you can control the frame rate here, set the lowest rate that provides enough detail for your sampling tasks. If you do not need colour and you can request monochrome here, you will save disk space (in compressed images) if you select monochrome.

If your video device includes audio support, there may also be audio setup controls here. These may replace or extend the controls in the **Audio Device Properties** dialog.

Audio Device Properties ...

This menu item appears in the Settings menu when you set an audio device. It opens a Properties dialog for your selected device. If the audio device has multiple inputs you can use this dialog to select the one to record. You can also set the volume level of the input. There are fields for stereo panning, treble and bass, loudness (bass boost) and to force monophonic sound from a stereo source.



With most common sound devices the majority of fields are grey and cannot be used. If the **Enable** field is not grey, it is important that you use the **Pin Line** field to display the input to record from and then check the **Enable** field to select that input. You should also make sure that the volume control (the vertical slider under the **Pin Line Input Mix** label) is not set to the bottom position, which is zero volume.

Video Capture...

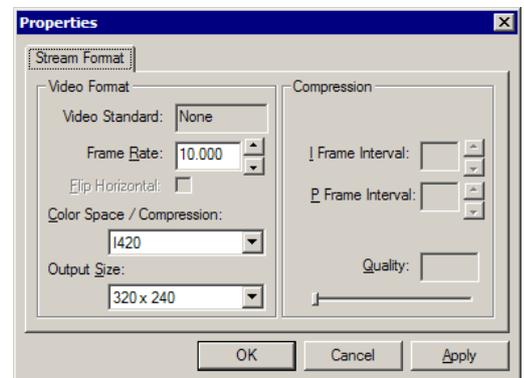
This menu item appears when you set a video device. It opens a dialog in which you can set the frame rate. It may also allow you to control the image size and format. This may duplicate controls in the Video Device Property page.

The **Frame Rate** sets the frames per second produced by the camera. You will get fewer frames per second than this if the system cannot keep up or if you limit rate with the **Set Slow Frame Rate** dialog or with the **Spike2**

`MMRate()` script command. The lower the frame rate, the smaller the data file. Use the lowest frame rate that gives you enough information for your task. A higher frame rate uses up processor time and disk space. It is likely that a rate of 10 or even 5 frames per second is all you need for many applications. The range of allowed frame rates depends on the camera. Some cameras have a single, fixed frame rate.

The **Output Size** is the image resolution expressed as horizontal x vertical pixels. Please use the smallest image size that is suitable for your needs. The size of the output file is proportional to the product of the horizontal and vertical sizes. Doubling the resolution creates a file that is four times the size.

The remaining settings depend on the device. Unless you have a good reason to change them we suggest you accept the initial values.

**Set Slow Frame Rate...**

This command opens a dialog in which you can set the maximum frame rate to capture and a slower rate to use when you do not need the rate set in the **Video Capture** dialog. The slower rate is also set by the **Spike2** `MMRate()` script command. We achieve these rates by throwing away frames captured by the camera.

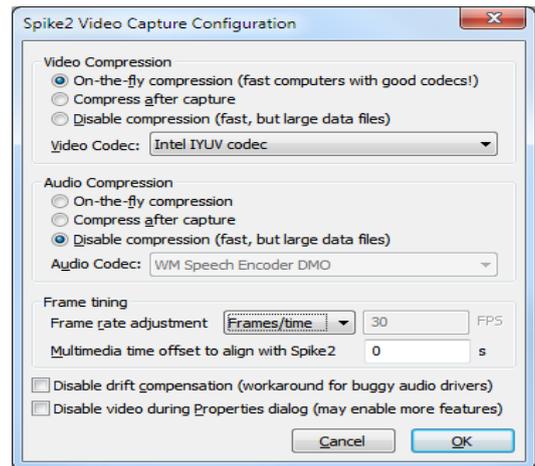
Use Slow Frame Rate

Use this command to swap between the slow frame rate and the maximum frame rate set in the **Set Slow Frame Rate** dialog.

Configuration... The configuration command opens a menu in which you can choose a data compression strategy and also how to adjust the data frame rate so the video stream stays synchronised to the times recorded in Spike2. It also has a check boxes for working around problems with some types of hardware.

Video and audio compression settings

Compressing your data can reduce file size considerably; reducing a file to less than 2% of its original size is possible especially if the video image is relatively constant. Given that a raw video rate is several MB a second, this is well worth having. The down side is that compressing data takes significant processing power. Ideally, you would like to compress the data in real time, but in practice this may not be possible. For both video and audio you can choose from:



On-the-fly compression For video, unless you have a very powerful computer or specialised hardware, this may be too slow to be useful unless you are running at low frame rates. If your computer cannot keep up it will drop frames from the video stream. Real time audio compression is usually possible. However, if you choose to compress the video after sampling, you may prefer to defer audio compression also.

Compress after capture If you select this option, the data is recorded raw to a temporary file. Make sure you have considered how large this file can get. Recording 640x480 images at 30 fps can easily use more than 1 GB a minute. When sampling has stopped and you choose to save the data file, the temporary file is added to a queue of files to be compressed. The compression task runs in the background, so that the program can continue capturing data. Spike2 will not be able to access the compressed files until the compression task has finished.

Disable compression This gives you the fastest sampling and access to your files the cost of the largest files.

Codecs You can choose to compress the video and audio data so that it occupies less disk space. This is done using plug in modules called codecs. A codec is in two parts: a coder and a decoder, hence codec. The coder changes the input signal in some way and the decoder reverses the process. Not all codecs are compressors; some change the format of the data to suit a particular purpose or to make the content secure. We are interested in codecs that can make the data significantly smaller and that do it fast enough to be useful.

Some codecs are lossless, that is, the decoded data exactly matches the encoded data; these tend to produce only small amounts of compression in real world data. Most codecs are lossy; the decoded result is not exactly the same as the original. However, they are designed to lose information in parts of the signal that we are less sensitive to, so often the result may look and sound close enough to the original not to matter. Different codecs are also designed for different purposes. For example, the Microsoft RLE video codec is designed to compress images with areas of identical colour, so it does a very good job on cartoons but makes a terrible mess of real-world images.

When you installed DirectX, a set of standard codecs was added to your system. It is likely that when you installed your video camera or web cam, the installation software added further codecs. The Audio codec and Video codec fields let you select codecs that

advertise themselves on your system as supporting audio or video. Not all of these are useful, or will even work with s2video.

The codecs listed below are those I found pre-installed on my quad core Windows 7-64 system in November 2011, plus DivX, which I installed as I knew of it. There are, no doubt, many other excellent codecs in existence, their omission from the lists is just an admission of my ignorance. This is not a scientific test but it give a flavour of what you can expect without doing any specialised codec set up to tune them for the data.

Video codecs To give you some ideas of which video codecs you could consider, the following table gives number of KB per second required to record of an image of me moving about in front of a web camera. The camera was set to produce 320x200 pixel images in colour at 30 frames per second. I converted the images on-the-fly; if the codex was slow, it dropped frames to keep up.

Codec	KB/sec	Ratio	Comments
None	6900	1	Uncompressed video
DivX	119	58	Good, fast (uses multiple CPU cores)
Microsoft Video 1	163	42	Blocky image
Cinepak Codec by Radius	230	30	Blocky image, dropped frames
Intel IYUV	2188	3.2	Good image
MJPEG Compressor	2584	2.7	Good image
DV Video Encode	3740	1.8	Good image

For 640x480 images, the data rate could be 4 times higher. For HD images (1920x1080), the data rates may be 27 times higher; you may need specialised hardware compressors.

Audio codecs There is also a wide range of audio codecs to choose from. Some are optimised for a particular task, such as compressing speech. The table shows the result of compressing me speaking the numbers 1 to 14 in approximately 11 seconds using stereo 16-bit sound at 44.1 kHz (which is the minimum setting allowed by my sound card).

Codec	KB/sec	Ratio	Comments
None	180	1	Uncompressed audio
PCM	180	1	This is a format converter
WM Speech Encoder DMO	8.3	22	Muffled but understandable
GSM 6.1	15	12	OK
WMAudio Encoder DMO	47	3.8	OK
IMA ADPCM	55	3.3	OK
Microsoft ADPCM	55	3.3	OK
CCITT A-Law	97	1.9	OK
CCITT u-Law	99	1.8	OK

The uncompressed file rate of 180 kB/s is 44.1 kHz times 4 bytes per sample plus overhead. You can reduce the audio data rate significantly by using 8-bit mono sound input (if your sound hardware allows this), which gets you down to 45 kB/s or so without a codec. If you can tolerate a lower sampling rate, this will also reduce the rate. Remember that you are trading off disk space, which mainly means the time taken to write to disk, against the time taken to encode the data. For example, if your sound hardware allows you to sample mono 8-bit data at 20 kHz and you are also sampling video, you may decide that the audio is the least of your problems and it is not worth the time to compress it.

If the exact synchronisation of the sound with other spike channels is important, record the sound as a Waveform channel. It is possible to switch recording of specific channels on and off to save disk space.

Disable drift compensation Time in a Spike2 data file is measured based on a crystal-controlled clock inside your 1401 interface. Time in a multimedia file is based on either the computer system clock or on a clock in the multimedia hardware. The 1401 interfaces use crystals that are accurate to 50 parts per million over a 0 to 70 degrees centigrade range, a worst case gain or loss of 4.3 seconds a day; most units will be better than that. Assuming that the clock in your computer is no more accurate than the 1401, and it could be much worse, there is a possibility of the multimedia recording and the 1401 disagreeing about time by as much as 8.6 seconds per 24 hours. To avoid problems with time drift, Spike2 keeps track of the time differences between 1401 time and computer time and notifies the `s2video` application of any drift. This drift can then be compensated for to make sure that the times in the multimedia file remain locked to those in the Spike2 data file.

Unfortunately, we have come across audio device drivers that crash when we try to adjust the time. Usually you can fix this by getting the most up-to-date driver for the audio device, but if this is not possible you can work around this problem by disabling the time drift compensation.

Do not confuse problems of time drift (which are usually small) with gross effects due to video devices claiming to run at one frame rate when they actually run at a different one. For example, I have a USB camera that claims to run at 30 frames per second, but actually runs at 31.8 frames per second. See the Configuration dialog to compensate for this kind of problem.

Frame timing AVI files hold a list of frames (which can be empty of data if the frame was dropped) at a constant time interval. The frame rate is the number of frames per second and this must be set accurately if Spike2 is to be able to scroll a video display to correspond with a time in a data file. When capture of a video file ends the frame rate is set in the AVI file and you can choose to control how the frame timing is recorded.

Frame rate adjustment There are three options: None, Frames/time and Fixed value.

- None Accept the frame rate generated by Windows, which will usually be the rate that the source hardware claims to generate.
- Frames/time Count the number of frames stored into the AVI file and divide this by the sample time in seconds. If the video really does run at a constant rate, this should generate a video file that lines up in Spike2.
- Fixed value Set the frame rate to a user-defined value. If you are triggering the frames with a clock that is known synchronised to Spike2, such as a clock generated by the output sequence Clock output, then you will know the exact clock period, so use this mode. We know of one triggerable camera that seems to guess a rate based on the image size when set to triggered mode. Setting the known rate works around this problem.

Multimedia time offset to align with Spike2 There will almost always be a time offset between the time that Spike2 starts running and the time that the camera records its first frame. This can be positive (if the first frame is late) or negative (if the system buffers up a few frames, ready for you to say go). If your device tends to have the same offset each time, you can set it here, and we save it in the AVI file. Spike2 will apply this offset when you review the file. The time is saved to millisecond precision.

This field is experimental and we store the information in a supposedly unused section of the AVI file header. AVI files not generated by `s2video` will always have this value set to 0. You can also adjust this value in Spike2 in the multimedia display window to line up the video exactly, assuming that the frame rate is correct.

How to check video timing If you set the following sequencer script:

```

set 1,1,0
go:    digout [00000000]
       delay s(5)-2
       digout [.....11]
       delay s(5)-2,go
    
```

This causes the front panel LEDs on Micro1401 and Power1401 devices to switch on and off every 5 seconds. If you connect digital output 0 to event 0 input on the front panel, then sample the input as a Level event channel, you can record the precise timing of the digital output changing state. If you point your video camera at the front panel, you can record a video that should show the front panel LEDs changing at exactly the times when the Level event channel changes. If the frame rate is exactly right, you are likely to see that the changes are shifted by a more or less constant time (remember that no times can be more accurate than the frame rate). If this delay stays constant over a range of video files, you can set this delay value in the dialog.

Working around problems Windows allows manufacturers to generate drivers for a huge range of devices. The interface specifications are complex and not all devices implement all aspects of the specifications perfectly. Because of this, we have added workarounds for problems that occasionally occur.

Disable drift compensation This is a workaround for some DirectShow audio device drivers that crash or generate errors when we try to match the multimedia time base to the Spike2 time. If you have problems of this nature, try checking this box (but the multimedia time and Spike2 time will drift apart). You can read more about drift compensation here. If you set Frames/time in the Frame rate adjustment there can sometimes be an advantage to checking this box as you are going to rescale the timebase using the frame rate.

Disable video (and audio) during Properties dialog We normally leave the multimedia capture running when you open the device properties dialogs. If you have Preview enabled, this allows you to see (and hear) the effect of any changes you make. However, some devices disable aspects of their Properties dialog if the device is being used. Check this box to stop running the device while the Properties dialogs are open.

View menu The view menu contains four items:



Toolbar

The toolbar holds short-cuts to preview the video image, to reduce the frame rate to the value set in the Settings menu Set Slow Frame Rate and to open the configuration menu. You can choose to show or hide it with this command.

Status bar You can choose to show or hide the status bar at the bottom of the window. The status bar gives you feedback about your menu selections and information about recording.

Always on top Tick this menu item to keep the s2video application on top of other windows.

Preview When a video camera is selected this shows or hides the video preview window.

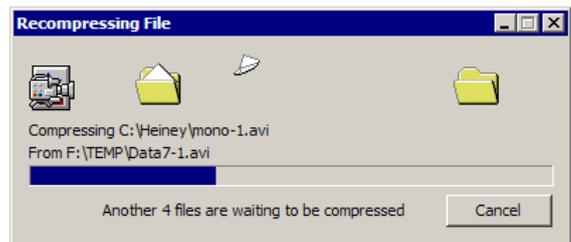
File menu This contains a single option, Exit, which will close the application unless you are recording or compressing data files.

Recording data

The `s2video` application acts as a slave to Spike2. If Spike2 is not running, you can configure the program, but it will not record data. Once a copy of Spike2 is located, `s2video` registers itself as a "listener" application. Spike2 will now tell it about all recording activities: opening a new temporary data file ready to record, starting recording, aborting or resetting sampling, stopping sampling, saving the temporary data file as a permanent file or throwing it away.

The `s2video` application does exactly the same tasks with its own multimedia file. There is one extra step required if you have chosen the **Compress after Capture** option in the configuration dialog. Instead of moving the temporary file to its final destination, it must be compressed. Compression can take a long time, and Spike2 might want to sample another file immediately. To avoid holding up Spike2, the temporary file is added to a list of files that need compression and the compression is run as a background task.

Compressing a file can take a long time, so while it happens `s2video` displays a dialog showing how far through the task the compressor has got for the current file. If there are multiple files waiting in the queue, the number of files left to process is also displayed.



In long recording sessions, most of the time you may need one video frame every few seconds, but there are short periods where you need the full frame rate. There is a Spike2 script command `MMRate()` that you can use to request a new frame rate. This command can only generate rates up to the maximum frame rate set for the camera.

When recording data under script control, you must use `FileSaveAs(name$, -1)` to save the data file and multimedia files. `FileSaveAs(name$, 0)` exports the data file and does not save any multimedia files.

If you get video stream timing problems, check that any frame rates set in the Video Device Properties match those in the Video Capture dialog. Beware of **Auto** options that may set unexpectedly high rates; it is usually better to set the rates manually, so that you know what they are. If the problems persist, set a lower frame rate or a smaller image size; the most common problem is that the system cannot process data fast enough.

Timing and AVI files

An AVI data file is a container for streams of data; in our case we are interested in audio and video data.

Video data is stored as frames, either as key frames that contain all the data needed to render them, or as frames that hold the difference from the previous frame in some way. In a standard AVI file, the first frame starts at time 0, and the frames occur at a fixed time interval. There is an index in the file that locates the position and holds the size of each frame. If frames were dropped during sampling, they still get index entries (to maintain the file timing), but the size of the data is usually 0. Video source devices announce the frame rates that they support, and in theory, they deliver frames at this rate. In practice, they deliver frames at a rate that is close to this rate, but can differ from it by several percent. When we also take into account the fact that the timebase used by Spike2 is based on a clock in the 1401 device, there is clearly a problem.

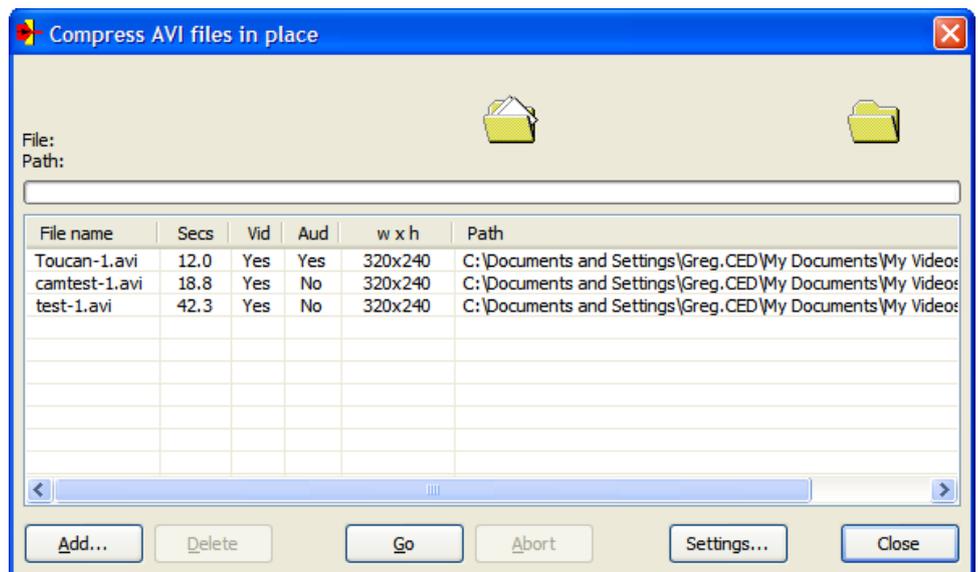
To help to correct time drifts, we adjust the clock used to time the multimedia to agree with the Spike2 clock during sampling. However, this doesn't do anything to correct the frame rate stored in the data file. If timing of video frames is a problem, there are options

in the Configuration menu to recalculate the frame rate based on number of frames and elapsed time. You can also set an initial time offset to account for fixed time offsets between the multimedia stream and the Spike2 stream.

If there is insufficient bandwidth in your system to write the multimedia stream to disk, data will be lost. In some cases, the system will be able to drop frames of data and still maintain the timing; in other cases this will make it impossible to line up the resulting AVI file with Spike2 data.

The avicomp application

The `avicomp` application can be used to compress `.avi` files generated by `s2video` that were not compressed during data capture. It compresses data files in place, to emulate what would have happened if you had chosen to compress them either during or after data capture. To run the application, double-click the file `avicomp.exe` that is in the same place as the Spike2 application.



The application main window is in three sections. The upper section holds the name and path of a file that is being compressed, with a progress bar that indicates how much of the file has been processed. The middle section holds a list of files that are to be processed and also displays the length of the file in seconds, if it holds video and audio and the size of any video image. The bottom section holds a list of buttons that you can use to control the program. The buttons are:

Add This button opens a standard file selection dialog in which you can select multiple `.avi` files. To be added to the list, each file must contain an AVI file (the name need not have the `.avi` extension, but all files created by `s2video` do have this extension). The files must also be readable and writable. You can always add a file, even when compressing. You cannot add a file that is already in the list.

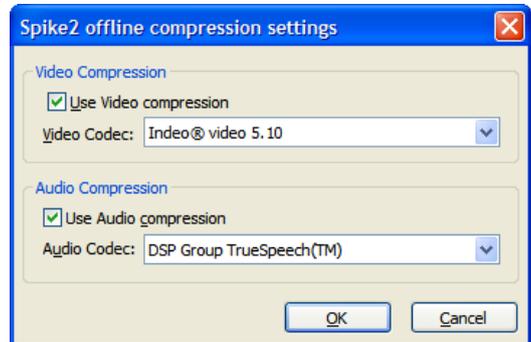
Delete This button is enabled when there is a selection in the list. The command removes the selected files from the list.

Go Start compressing the files. Each file is removed from the list as it is compressed and the file name, path and compression progress are displayed in the upper part of the display. Files are compressed to a temporary file in the same folder, then if the compression is

successful, the original file is deleted and the temporary file is renamed to the original name. If you want to preserve the original file, you must save a copy first.

Abort Click this button to stop compressing the current file. Any temporary file is deleted and the file will be put back into the list. The compression process will stop.

Settings... This command opens a dialog where you can choose the video and audio compression filters to apply to your file. You can change settings, even while processing files. The settings used will be those in effect when each file starts to be processed. You must enable some form of compression or the OK button is disabled. The choice of codec is the same as for s2video.



Close This button is enabled when you are not compressing any data. It saves the list of file names in the window (up to a maximum of 100 files) in the registry, then closes the application. The list of files will be reloaded the next time you run the avicomp application.

Index

—Symbols—

- : in Ternary operator for dialogs
 - Colon, 3-2
- : in times for dialogs
 - Colon, 3-2
- ? in Ternary operator for dialogs, 3-2

—1—

- 10 MHz clock, 4-13
- 1401
 - Device driver, 14-2
 - Monitor revision, 14-2
 - Select unit, 3-6
- 1401 type, 4-12
- 1401-18, 4-7, 19-1
 - Sample menu option, 12-4
- 1902
 - AC coupling, 18-2
 - Interactive support, 18-1
 - Trigger inputs, 18-2

—A—

- Abort sampling, 4-28
- Abs Max value
 - between cursors, 11-7
- ABS output instruction, 5-38
- Abs () virtual channel function, 9-32
- Active cursors, 11-3
 - Step left/right, 2-3, 3-3
- Add items to memory buffer, 9-27
- ADD output instruction, 5-39
- ADDAC output instruction, 5-27
- ADDACn output instruction, 5-28
- ADDI output instruction, 5-38
- Aliasing of waveform data, 4-4
- All pass filter, 17-5
- All stop filter, 17-5
- Analysing data, 2-8
- Analysis menu, 9-1
 - Delete channel, 9-35
 - Fit data, 9-21
 - Measure to a channel, 9-20
 - Measurements, 9-14
 - Memory buffer, 9-26
 - New Result View, 9-1
 - Process settings, 9-11
 - Save channel, 9-35
- AND, ANDI output instructions, 5-40
- ANGLE output instruction, 5-32
- Arbitrary waveform output, 4-20
- Area
 - between cursors, 11-7
 - Modulus between cursors, 11-7
 - Modulus in trend plot, 9-19
 - trend plot, 9-18
 - under curve between cursors, 11-7

under curve in trend plot, 9-19

- Area under curve, 2-6
- Arrange icons, 10-2
- ASz () sequencer expression, 5-19
- ATan () virtual channel function, 9-32
- Audio codecs, 21-5
- Audio recording, 21-1
- Auto Format, 7-9
- Auto-correlation of events, 9-4
- Auto-correlation of waveforms, 9-11
- Automatic file naming, 4-24
 - Date and time specifiers, 4-24
- Automatic file saving, 4-24
- Automatic scrolling, 8-21
- Automatic update of result view, 9-13
- Auxiliary values, 8-14
- Auxiliary variables, 8-14
- Average. *see* Mean
- Average of waveform, 9-7
 - Number of sweeps, 8-4
- avi file extension, 6-1, 8-16
- avicom application, 21-9
- Axis controls, 2-4, **8-1**
 - Show and Hide, 2-5, **8-3**

—B—

- Band pass filter, 17-5
- Band stop filter, 17-5
- Duplicate channels, 15-14
- beats per minute, 8-10
- BEQ output instruction, 5-37
- Bessel filter, 17-8
- BGE output instruction, 5-37
- BGT output instruction, 5-37
- Big file, 6-5
 - Explanation, 4-2
- Bitmap output
 - To clipboard, 7-1
 - To file, 6-5
- Black and white displays, 8-20
- BLE output instruction, 5-37
- BLT output instruction, 5-37
- BNE output instruction, 5-37
- BNZERO output instruction, 5-26
- Bookmarks, 7-8
 - Set on found text, 7-7
- bpm display mode, 8-10
- BRAND output instruction, 5-43
- BRANDV output instruction, 5-43
- Break out of cursor calculations, 11-7
- Breaking out of processing, 9-12
- Breaths per minute, 8-10
- Buffer overflow, 4-29
- Butterworth filter, 17-8
- Bxx output instructions, 5-37
- BZERO output instruction, 5-26

—C—

- C0-C9 cursor shortcut in dialogs, 3-2
- Calibrate waveform
 - Dialog, 9-35
 - methods, 9-36
- Call tips
 - Display of, 7-13
 - Style for, 7-12, 7-14
- CALL, CALLV output instructions, 5-36
- CANGLE output instruction, 5-32
- Cascade windows, 10-1
- Case sensitivity
 - In searches, 7-7
 - Output sequencer, 5-18
- CED 1902, 18-1
- CED Software help desk, 14-1
- CEDCOND.LOG file, 7-17, 18-1
- cfbx file extension, 6-1
- Ch (n) virtual channel function, 9-30
- CHAN output instruction, 5-42
- Channel arithmetic, 9-28
- Channel display
 - Cubic spline, 8-12
 - Dots, **8-10**
 - Draw mode colour, 8-18
 - Instantaneous frequency, 8-11
 - Lines, **8-10**
 - Mean frequency, 8-10
 - Order, 7-16
 - Overdraw WaveMark mode, 8-11
 - Raster mode, 8-11
 - Rate, 8-11
 - Sonogram mode, 8-11
 - State mode, 8-13
 - Text mode, 8-13
 - Waveform mode, 8-12
 - WaveMark mode, 8-12
- Channel duplication, 9-34
 - Titles, 8-4
- Channel lists, 3-1
- Channel parameters, 4-3
- Channel process
 - Arbitrary waveform output, 12-4
 - Calibration, 9-36
 - Dialog, 9-38
- Channel range, 3-1
- Channel selection, 2-1
- Chebyshev type 1 filter, 17-8
- Chebyshev type 2 filter, 17-8
- Clear template, 15-5, 15-6
- Clear text or result view, 7-1
- Clipboard
 - Copy result view as text, 7-2
 - Copy spike shapes, 15-12
 - Copy time view as text, 7-3
 - Copy to, 7-1
 - Copy XY view as text, 7-3

- Cut text to, 7-1
- Paste text, 7-1
- Clock output during sampling, 5-3
- Clock rate for output sequencer, 5-2
- Clock tick size, 2-9
- Clock**, front panel connection, 5-3
- Close all associated windows, 6-4, 10-1
- Close document, 6-4
- CLRC output instruction, 5-34
- Clustering
 - Apply changes, 16-11
 - Clustering dialog, 16-8
 - Copy as bitmap, 16-13
 - Copy cluster values, 16-13
 - Copy data as text, 16-13
 - Delete all online spikes, 16-12
 - Display principal components, 16-12
 - Ellipse control, 16-10
 - Ellipse radius, 16-14
 - from correlations, 16-6
 - from errors, 16-7
 - from measurements, 16-4
 - Getting started, 16-20
 - Interval histogram, 16-12
 - INTH settings, 16-13
 - Introduction, 16-1
 - Jiggle settings, 16-16
 - K Means algorithm, 16-21
 - K Means dialog, 16-17
 - Mahalanobis distance, 16-23
 - Match to classes, 16-19
 - Menu commands, 16-11
 - Normal Mixtures algorithm, 16-23
 - Normal Mixtures dialog, 16-18
 - Online update, 16-11
 - Principal Component Analysis, 16-2
 - Reanalyse data, 16-11
 - Restore class codes, 16-11
 - Select Principal Components, 16-12
 - Set codes, 16-16
 - Time range control, 16-19
 - Tracking changes over time, 16-19
 - View along axis, 16-16
 - View settings, 16-14
- Codec in multimedia recording, 21-4
- Coefficients of filters, 17-9
- COFF output instruction, 5-33
- Collision Analysis Mode, 15-14, 15-15
- Colour dialog, 8-18
 - Save colours in registry, 7-19
- Command line options, 3-6
- Comment
 - Data channel, 4-3
 - Display in Windows Explorer, 3-6
 - File comment, 8-4
 - File comment auto-prompt, 4-24
 - Output sequencer, 5-18
- Compatibility of Spike2 versions, 1-3
- Settings for, 7-19
- Compile output sequence, 5-14
- Compile script, 13-1
- Conditioner
 - CEDCOND.LOG file, 7-17, 18-1
 - Sample menu, 12-4
 - Serial port, 7-17, 18-1
- Configuration files, 6-1
 - Contents, 4-30
 - Load and run, Sample Bar, 12-1
 - Load and save, 6-6
- Connections
 - Digital i/o for sequencer, 5-4
 - Event 0 and 1, 7-16
 - Event discriminator, 19-4
 - Event inputs, 4-5
 - Event inputs, 4-5
 - Power1401 DACs 2 to 5, 5-27
 - Sample and Play wave trigger, 7-16
 - Sequencer clock output, 5-3
 - Signal conditioner, 7-17, 18-1
 - Template match and DIGLOW, 15-2
 - Waveform channels, 4-4
 - Waveform output, 5-27
- Continuous sampling mode, 4-18
- Convert
 - DOS script, 13-2
 - Event to waveform, 9-28
 - foreign file format, 6-2
 - RealMark to waveform, 9-28
- Copy
 - Cluster information as text, 16-13
 - Clustering data as bitmap, 16-13
 - Clustering data as text, 16-13
 - Result view as text, 7-2
 - Result, Time and XY as images, 7-1
 - spike shape templates, 15-12
 - Spreadsheet format, 7-2
 - Text, 7-1
 - Time view as text, 7-3
 - XY view as text, 7-3
- Copy channels (sampling), 4-13
- Correlation of events, 9-4
 - Number of sweeps, 8-4
- Correlation of waveforms, 9-11
- COS () virtual channel function, 9-32
- Cosine wave output, **5-29**
- Count of events, 2-6, **11-8**
- CPHASE output instruction, 5-32
- Crash recovery, 20-1
- CRATE output instruction, 5-31
- CRATEW output instruction, 5-31
- Create new channel, 9-28
- Create new memory buffer, 9-26
- Create TextMark, 12-1
- CRINC output instruction, 5-34
- CRINCW output instruction, 5-34
- CSZ output instruction, 5-30
- CSZINC output instruction, 5-30
- Cub () Virtual channel function, 9-32
- Cubic spline display mode, 8-12, 8-13
- Cursor 0, 11-3
- Cursor 0 stepping, 9-14
- Cursor button, 2-2
- Cursor dialog, 2-6
- Cursor labelling styles, 2-2, 11-2, 11-10
- Cursor menu, 11-1
 - Active mode, 11-3
 - Cursor regions, 11-7
 - Delete, 11-1
 - Delete Horizontal, 11-9
 - Display all, 11-1
 - Display all Horizontal, 11-9
 - Move To, 11-1
 - Move To Horizontal, 11-9
 - New cursor, 11-1
 - New horizontal, 11-9
 - Renumber cursors, 11-2
 - Renumber horizontal cursors, 11-10
- Cursor Regions, 11-7
- Cursors
 - Active, 11-3
 - Adding, 2-2, **11-1**
 - Adding horizontal, **11-9**
 - Break out of long calculations, 11-7
 - Copy value to Log view, 11-6, 11-8
 - Delete cursor, 11-1, 11-9
 - Display all, 11-1, 11-9
 - In result view, 2-10
 - Label style, 11-2, 11-10
 - Move window to cursor, 11-1
 - Move y axis to cursor, 11-9
 - Print cursor values, 11-6, 11-8
 - Renumber, 11-2
 - Renumber Horizontal, 11-10
 - Static, 11-4
 - Valid and invalid, 11-3
 - Value at, 2-6, **11-6**
 - Values between, 2-6, **11-7**
- Curve fitting, 9-21
 - Testing the fit, 9-24
- Cut text, 7-1
- CWAIT output instruction, 5-33
- CWCLR output instruction, 5-34
- CyberAmp Interactive support, 18-1

—D—

- DAC connections for Power1401, 5-27
- DAC output during sampling, 5-1
 - Arbitrary waveforms, 4-20
- DAC output instruction, 5-27
- DAC output offline, 12-3
- DAC outputs
 - Reset, 4-26

- DACn output instruction, 5-28
 - Damaged files, recovery, 20-1
 - DANGLE output instruction, 5-32
 - Data channels, 2-1
 - Data points cursor mode, 11-5
 - Data view short cut keys, 3-3
 - Date in automatic file name, 4-24
 - dB scale, 17-10
 - DBNZ output instruction, 5-35
 - DBNZn output instruction, 5-36
 - DC remove wave, 9-39
 - Debounce event inputs, 4-5
 - Debounce marker input, 4-6
 - Debug bar, 13-1
 - Decibel scale, 17-10
 - default.s2c configuration, 8-20
 - DEFAULT.S2C configuration, 6-6
 - DELAY output instruction, 5-35
 - Delete data channel, 9-35
 - Delete memory buffer items, 9-27
 - Delete selected text, 7-1
 - DF () virtual channel function, 9-30
 - Dialog expressions, 3-1
 - DIBEQ output instruction, 5-25
 - Clash with template match, 15-2
 - DIBNE output instruction, 5-25
 - Dichotic notch, 9-15
 - Differentiate wave, 9-39
 - Differentiator, 17-11
 - Differentiator filter, 17-6
 - DIGIN output instruction, 5-26
 - Digital filter
 - Dialog, 17-2
 - FIR filter dialog, 17-4
 - FIR filters, 17-1
 - FIRMake filter types, 17-11
 - IIR filter dialog, 17-6
 - IIR filters, 17-1
 - Digital inputs, 4-5
 - Conflict with Marker data, 5-3
 - Connections, 5-4
 - marker connections, 4-6
 - test bits, 5-25
 - test saved bits, 5-25
 - TTL levels, 4-5
 - wait for bit pattern, 5-26
 - With 1401-18, 4-7, 19-4
 - Digital marker, 4-6
 - Connections, 5-4
 - Link to output, 4-7
 - With 1401-18, 19-4
 - Digital outputs
 - Connections, 5-4
 - DIGLOW, 5-24
 - DIGOUT, 5-24
 - NDR and NDRL, 5-4
 - Reset, 4-26
 - Sequencer update, 5-2
 - Template match, 15-2
 - WaveMark template match, 15-2
 - DIGLOW output instruction, 5-24
 - Template match clash, 15-2
 - DIGOUT output instruction, 5-24
 - Directory for new data files, 7-17
 - DISBEQ output instruction, 5-25
 - DISBNE output instruction, 5-25
 - Discriminator, 19-1
 - Sample menu option, 12-4
 - Display trigger, 8-6
 - DIV output instruction, 5-39
 - Time penalty, 5-2
 - DOFF output instruction, 5-33
 - DOS scripts, 13-2
 - DOS version of Spike2, 1-3
 - Dot size, cluster window, 16-14
 - Dots draw mode, 2-5, **8-10**
 - Down sample wave, 9-39
 - DPHASE output instruction, 5-32
 - Drag and drop text, 3-5
 - DRange () sequencer expression, 5-19
 - DRATE output instruction, 5-31
 - DRATEW output instruction, 5-31
 - Drawing modes
 - Result view, 8-13
 - Time view, 8-10
 - DRINC output instruction, 5-34
 - DRINCW output instruction, 5-34
 - DSZ output instruction, 5-30
 - DSZINC output instruction, 5-30
 - Dummy channels, 4-13
 - Duplicate channels, 9-34
 - Titles, 8-4
 - DWAIT output instruction, 5-33
 - DWCLR output instruction, 5-34
- E—**
- EC () virtual channel function, 9-31
 - Edit menu, 7-1
 - Auto Format, 7-9
 - Clear, 7-1
 - Copy, 7-1
 - Copy as text, 7-2, 7-3
 - Copy Spreadsheet, 7-2
 - Cut text, 7-1
 - Delete selection, 7-1
 - Find text, 7-7
 - Paste, 7-1
 - Preferences, 7-11
 - Replace, 7-8
 - Select All, 7-6
 - Undo, 7-1
 - Edit text
 - Bookmarks, 3-5, 7-8
 - Cut, copy and paste, 3-4
 - Drag and drop, 3-5
 - Find, 3-5, 7-8
 - Indent and outdent, 3-5
 - Rectangular select, 3-6
 - Regular expressions, 7-7
 - Search, 7-7
 - Short cut keys, 3-3
 - Text caret control, 3-3
 - Wildcard searches, 7-7
 - Edit TextMark, 12-1
 - Edit toolbar, 7-8
 - Edit WaveMark data, 15-13
 - Eg () virtual channel function, 9-31
 - Ellipse control in clustering, 16-10
 - Email support, 6-10
 - End of line characters, fixing, 7-1
 - Errors, output sequencer compiler, 5-46
 - ES () virtual channel function, 9-31
 - Escape from drawing, 8-14
 - Et () virtual channel function, 9-31
 - Evaluate a script line, 13-1
 - Event 0 and 1 connections
 - Digital input pins, 4-5
 - Micro1401 and Power1401, 7-16
 - Event 3 sampling trigger, 4-6, 4-28
 - Event count, 2-6, **11-8**
 - Event data
 - Connections, 4-5
 - Convert to waveform, 9-28
 - Copy As Text, Export As format, 7-5
 - Find next, last, 2-3, 3-3, 8-2
 - Measurements and draw mode, 7-19
 - Event discriminator, 4-7, 19-1
 - Connections, 19-4
 - Electrical information, 19-4
 - Mode, 19-2
 - Monitor channel, 19-3
 - Route of signals, 19-1
 - Sample menu option, 12-4
 - Event display
 - Dots, 2-5
 - Lines, 2-5
 - Exit, 6-9
 - Export data
 - As bitmap file, 6-5
 - As spike2 file, 6-5
 - As spreadsheet text, 6-5
 - As text file, 6-5
 - As Windows Metafile, 6-5, 7-15
 - TextMark, 12-2
 - To clipboard, 7-2, 7-3
 - Expressions in dialogs, 3-1
 - External exporter, 6-4
 - Extreme value
 - trend plot, 9-19

—F—

FFT (Fast Fourier Transform), 9-8
 File comments, 8-4
 Automatic prompt, 4-24
 Display in Windows Explorer, 3-6
 File format converters, 6-2
 File icons, 1-2
 File menu, 6-1
 Close, 6-4
 Exit, 6-9
 Export, 6-4
 Global resources, 6-2
 Import, 6-2
 New File, 6-1
 Open, 6-2
 Page Setup, 6-6
 Print (Visible, Selection), 6-8
 Print screen, 6-9
 Resource files, 6-3
 Revert To Saved, 6-4
 Save and Save As, 6-4
 Send Mail, 6-10
 File name extensions, 6-1
 File size limits, 4-24
 Fill gaps, 9-40
`Filterbank.cfb` filter bank file, 17-3
 Filter bank, 17-3
 Filter marker data, 9-40
 Find next/last keyboard command, 3-3
 Find Short Interval, 16-15
 Find text, 7-7
 FIR filter, 17-2
 Attenuation and ripple, 17-11
 Coefficients, 17-9
 Differentiator example, 17-16
 Frequencies, 17-9
 Frequency bands, 17-10
 Maximum useful attenuation, 17-10
 Multiband example, 17-15
 Number of coefficients, 17-13
 Nyquist frequency, 17-14
 Overview, 17-1
 Pink noise from white noise, 17-12
 Ripple in bands, 17-11
 Slope for differentiator, 17-12
 Technical details, 17-9
 Transition region, 17-10
 Weighting, 17-10
 FIRMake()
 Discussion, 17-9
 filter types, 17-11
 Fit data, 9-21
 Flush sampled data to disk, 4-25
 Folding, 7-13
 Font selection, 8-17
 Format of Spike2 data files, 1-3

—G—

Gated update of result view, 9-13
 Global resource files, 6-2
 Gradient of line, 2-6, **11-7**
 Graphical sequence editor, 5-5
 Grid
 Set colour, 8-18
 Show and hide, 8-3
 Group channels, 8-2

—H—

HALT output instruction, 5-37
 Hamming window, 8-12
 Hanning window, 8-12
 Hardware required for Spike2, 1-3
 Header and footer, 6-7
 Print screen, 6-9
 Help, 2-1, **14-1**
 Help desk, 14-1
 Hexadecimal marker codes, 4-7
 Hide cursor 0, 11-3
 Hide scroll bar, 8-3
 Hide window, 10-1
 High pass filter, 17-5
 High pass filter example, 17-14
 High-pass filter wave, 9-39
 Hilbert transformer, 17-12, 17-18
 Hold triggered display, 8-6
`Hwr()` virtual channel function, 9-32
`Hz()` sequencer expression, 5-19

—I—

Icons for files, 1-2
 Icons, arrange, 10-2
 Ideal waveform sampling rate, 4-4
`If()` virtual channel function, 9-30
`Ifc()` virtual channel function, 9-30
 IIR filter
 Details, 17-6
 Filter model, 17-8
 Filter order, 17-7
 Filter type, 17-7
 Notch filter, 17-8
 Overview, 17-1
 Import channel into memory buffer, 9-28
`import` folder, 6-2
 Import foreign data file, 6-2
 Impulse response, 17-9
 Include files, 5-22
 Info on a result view, 8-4
 Installing Spike2, 1-4
 Instantaneous frequency, 8-11
 Instructions for output sequencer, 5-17
 Interpolate wave, 9-39
 Interrupt drawing, 8-14
 Interrupting cursor window, 11-7

Interval histogram, 2-8, **9-2**
 Linked to clustering, 16-12, 16-13
 Number of intervals, 8-4
 INTH. *see* *Interval histogram*
 Invalid cursors, 11-3

—J—

J3 clustering measure, 16-22
 JUMP output instruction, 5-37
 Jump to event, 8-2

—K—

K Means
 Algorithm, 16-21
 Dialog, 16-17
 Kaiser window, 8-12
 Key for XY view, 8-5
 Keyboard markers, 4-6
 Special code FF, 4-29
 Special codes, 4-18, 4-28
 Keyboard sequencer control, 5-1, 5-18

—L—

Label for cursor, 2-2, **11-2**, **11-10**
`LAST.S2C` configuration, 6-6
`LD1RAN` output instruction, 5-44
`LDCNTn` output instruction, 5-36
 Least squares fitting, 2-6
 Level crossing, 9-28
 Level event data draw mode, 8-10
 Licence information, 1-2
 Limits on sample time and file size, 4-24
 Line draw mode for events, 2-5, **8-10**
 Line style in XY views, 8-14
 Line thickness, 8-22
 Line thickness for printing, 7-15
 Load and Run a script, 13-1
 on startup, 3-6
 Load templates, 15-18
 Lock a template, 15-5
 Lock y axes, 8-2
 Log view line limit, 7-11
 Logarithmic axes
 X Axis, 8-1
 Y Axis, 8-3
 Logarithmic scale, 17-10
 Low pass differentiator filter, 17-6
 Low pass filter, 17-5
 Low pass filter example, 17-12

—M—

Macintosh
 68000 version, 1-3
 PowerPC native mode, 1-3
 Magnify pointer, 2-3
 Mahalanobis distance, 16-23

- Manual update of result view, 9-13
- Mark divider, 4-13
- MARK output instruction, 4-7, 5-42
- Marker data
 - Conflict with sequencer, 5-3
 - Copy As Text, Export As format, 7-5
 - Filter markers, 9-40
 - Set codes, 9-41
 - State display mode, 8-13
- Marker filter, 9-40
 - in Spike shape dialogs, 15-12
- Match channel, 9-40
- MATLAB export, 6-4
- Maximum possible run time, 4-12
- Maximum sustained event rate, 4-5, 4-6, 4-9
- Maximum total sampling rate, 4-3
- Maximum value
 - between cursors, 11-7
 - trend plot, 9-19
- Maximum, minimum cursor modes, 11-4
- Mean event rate, 11-8
- Mean frequency, 8-10
- Mean value
 - Between cursors, 2-6, **11-7**
 - trend plot, 9-18
- Measurements
 - Cursor positions, 11-10
 - Tabulated output of, 7-3
 - To a data channel, 9-20
 - To an XY view, 9-14
 - Types, 9-17
- Median filter, 9-40
- Memory
 - Limit on histogram size, 2-9
- Memory buffer, **9-26**
 - Add items, 9-27
 - Create new channel, 9-26
 - Delete items, 9-27
 - Write to file, 9-28
- Menus
 - Analysis, 9-1
 - Cursor, 11-1
 - Edit, 7-1
 - File, 6-1
 - Help, 14-1
 - Sample, 12-1
 - View, 8-1
 - Windows, 10-1
- Metafile image export, 15-12
- Metafile output
 - To clipboard, 7-1
 - To file, 6-5
- Metafile output resolution, 7-15
- MF () virtual channel function, 9-30
- Microseconds per time unit, 4-12, 4-15
- Minimum value
 - between cursors, 11-7
- trend plot, 9-19
- MOV output instruction, 5-38
- MOVI output instruction, 5-38
- MOVRND output instruction, 5-44
- ms as time scaler, 3-2
- ms () sequencer expression, 5-18
- mSTick () sequencer expression, 5-18
- MUL and MULTI output instructions, 5-39
- Multiband filter, 17-11
- Multiband with 3 dB/octave cut, 17-12
- Multimedia files, 8-16
- Multimedia recording, 21-1
 - avicom offline compressor, 21-9
 - Disable drift compensation, 21-6
 - Getting started, 21-2
 - Recording data, 21-8
 - Set slow frame rate, 21-3
 - System requirements, 21-1
 - Use slow frame rate, 21-3
 - Video Capture dialog, 21-3
- Multiple 1401s, 3-6
- N—
- NDR NDRL digital output signals, 5-4
- NEG output instruction, 5-38
- New document, **4-27**, 6-1
 - Temporary folder for data files, 7-17
- New file from existing file, 6-4
- New result view, 9-1
- Next data point search, 11-5
- NOP output instruction, 5-37
- Normal Mixtures dialog, 16-18
- Notch filter, 17-8
- Numeric expressions in dialogs, 3-1
- O—
- Off-line templates, 15-17, 15-18, 15-19
- OFFSET output instruction, 5-33
- One and a half high pass filter, 17-6
- One and a half low pass filter, 17-5
- On-line templates, 15-17, 15-18, 15-19
- Opening
 - Files from command line, 3-6
 - New document, **4-27**, 6-1
 - Old document, 6-2
- Operators
 - In dialogs, 3-2
- Optimise Y axis, 8-2
- Options for XY view, 8-5
- OR, ORI output instructions, 5-40
- Order of channels, 7-16
- Oscilloscope style trigger, 8-6
- Out, front panel connection, 5-3
- Output Reset, 4-26
- Output sequencer, 5-1
 - Access to data capture, 5-42
- Add constant to variable, 5-38
- Arbitrary waveform output control, 5-11, 5-44
- Bitwise logic instructions, 5-40
- Calculate variable values, 5-20, 5-24, 5-25, 5-27, 5-28, 5-34
- Change during sampling, 5-15, 5-23
- Clock rate, 5-2, 5-5
- Compare variable, 5-37
- Compatibility with old versions, 5-46
- Compile sequence, 5-14
- Compiler errors, 5-46
- Constants, 5-20
- Control panel, 5-1
- Copy variable, 5-38
- DAC outputs, 5-10, 5-27
- DAC scaling, 5-6, 5-21
- Disable interactive jumps, 5-1, 5-5
- Divide variable, 5-39
- Dragging and duplicating items, 5-9
- Expressions, 5-18
- Format text, 5-14
- Format with step numbers, 5-14
- Get current sample time, 5-42
- Graphical editing, 5-8
- Graphical editor setup, 5-5
- Graphical palette, 5-9
- Instruction format, 5-18
- Instruction reference, 5-24
- Instructions, 5-17
- milliseconds per step, 5-2, 5-5, 5-20
- Multiply variables, 5-39
- Negate variable, 5-38
- Randomisation, 5-12, 5-13, 5-14
- Randomisation, 5-43
- Reciprocal of variable, 5-39
- SCLK directive, 5-20
- SDAC directive, 5-20
- SET directive, 5-20
- Set file to use, 4-19
- Set variable value, 5-38
- TABDAT directive, 5-21
- Table of values, 5-21
- TABSZ directive, 5-21
- Timing faults, 5-9
- Variable sum, difference, 5-13, 5-39
- Variables, 5-19
- Overdraw data
 - based on trigger times, 8-6
- Overdraw WM display mode, 8-11
- P—
- Paged display on-line, 8-6
- Pass band, 17-10, 17-11
- Paste text, 7-1
- Peak and trough
 - between cursors, 11-8

trend plot, 9-19

Peak search, 9-28
 cursor mode, 11-4

Peak to peak
 between cursors, 11-8
 trend plot, 9-19

Pen Width, 8-22

per minute rates, 8-10

Peri-stimulus time histogram, 9-3
 Number of sweeps, 8-4

Phase histogram, 9-6
 Number of cycles, 8-4

PHASE output instruction, 5-32

Play waveform, 12-3

Play waveform toolbar, 4-20

pls file extension, 6-1

Point style in XY views, 8-14

POLY() virtual channel function, 9-32

Port for sampling, 4-3

Power spectrum, 9-8
 As waveform, 9-28

Power1401 gain option, 18-1

Preferences, 7-11

Preferences folder, 4-30

Principal Component Analysis, 16-2
 Mathematics, 16-3

Printing, 6-6
 Line thickness, 7-15
 Preview printed output, 6-8
 Spike shape templates, 15-12

Process dialog, 2-9, **9-12**
 for New file, 4-28, 9-13

Process settings, 9-11

Prompt to save result/XY views, 7-11

PSTH. *see* Peri-stimulus time histogram

Pulse outputs, 5-10

PW() virtual channel function, 9-30

—Q—

Quiet startup, 3-6

—R—

RAMP output instruction, 5-29

Randomisation in output sequencer, 5-43

Raster display
 Drawing modes, 8-14
 Event correlation, 9-5
 PSTH, 9-3

Raster event display mode, 8-11

Raster sweep values, 9-4, 9-5, 9-6

Rate drawing mode, 8-11
 Setting interactively, 8-10

RATE output instruction, 5-31

RATEW output instruction, 5-31

Read only. *See* Read-only

Read-only files, 6-10

RealMark data, 4-9

Convert to waveform, 9-28

Copy As Text format, 7-6

Export As format, 7-6

RealWave data, 4-9
 Export as Adc data, 6-5
 Rescale, 8-5

RECIP output instruction, 5-39
 Time penalty, 5-2

Reclassify WaveMark data, 15-13

Record a script, 13-1

Recover data files, 20-1

Rectify waveform
 Channel process, 9-38

Regions dialog, 2-6

Relative measurements, 2-6, **11-6**

Remove Spike2, 1-4

Renumber cursors, 11-2

Renumber horizontal cursors, 11-10

Replace matched text, 7-8

Replay waveforms, 12-3

Repolarisation cursor mode, 11-5

REPORT output instruction, 4-7, 5-42

Rerun an existing file, 8-21
 linked to play waveform, 12-3

Resample wave, 9-39

Reset sampling, 4-28

Resolution in time, 4-12

Resonator filter, 17-8

Resource files, 6-1
 Apply and save, 6-3

Result files (.srf), 6-1

Result view, 2-8
 Clear, 7-1
 Copy as text, 7-2
 Creating a new view, 9-1
 Cursors, 2-10
 Draw modes, 2-10, 8-13
 Mean, area, sum and slope, 11-7
 Number of sweeps or items, 8-4
 Prompt to save unsaved view, 7-11
 Rasters, 8-14
 Update mode, 9-13

RETURN output instruction, 5-36

Revert text document to last saved, 6-4

RINC output instruction, 5-34

RINCW output instruction, 5-34

Rm() virtual channel function, 9-30

Rmc() virtual channel function, 9-30

RMS Amplitude
 between cursors, 11-8
 Calibration method, 9-37
 Channel process option, 9-40
 Trend plot, 9-19

ROM revisions in 1401, 20-5

Rotating the cluster window, 16-9

rpm, 8-10

Run script, 13-1

command line, 3-6
 from Script Bar, 13-2
 startup.s2s, 3-6

—S—

s() sequencer expression, 5-18

s2cx Configuration file extension, 6-1

s2rx Resource file extension, 6-1

s2s script file extension, 6-1

s2video application, 8-17, 21-1

Sample Bar, **12-1**
 Label and comment, 4-25

Sample control toolbar, 4-28

Sample interval, 4-4

Sample menu, 12-1
 Create a TextMark, 12-1
 Discriminator command, 19-1
 Discriminator configuration, 12-4
 Sample Bar, 12-1
 Sampling configuration, 12-1
 Sequencer controls, 12-4
 Signal conditioner setup, 12-4
 Waveform output, 12-3

Sample rate for waveform data, 4-4

Sample Status bar, 4-29

Sampling
 Controls during, **4-28**, 12-4
 Diagnosing problems, 20-5
 Digital outputs, 5-1
 File size limit, 4-24
 Maximum rate, 4-3
 Mode, 4-17
 Multimedia data, 21-8
 Multiple files, 4-24
 Output during, 5-1
 Time limit, 4-24
 Triggered start, 4-28

Sampling configuration, **4-1**, 4-30, 12-1
 Add a new channel, 4-2
 Automation tab, 4-24
 Channel types, 4-3
 Channels tab, 4-1
 Contents of file, 4-30
 Loading and saving, 6-6
 Mode tab, 4-17
 Play waveform tab, 4-20, 4-22
 Resolution tab, 4-12
 Sequencer tab, 4-19
 Spike2 Last file, 4-30
 Wavemark data, 15-2

Sampling Notes, 12-2

Save data channel, 9-35

Save files, 6-4
 Automatic name generation, 4-24
 Automatic save of script, 7-11
 Automatically after sampling, 4-24
 during sampling, 4-25

- SCLK directive, 5-20
Screen dump, 6-9
Script
 Automatic save if modified, 7-11
Script Bar, **13-2**
Script menu, 13-1
 Compile script, 13-1
 Convert DOS script, 13-2
 Debug bar, 13-1
 Evaluate, 13-1
 Run script, 13-1
 Script Bar, 13-2
 Turn recording on and off, 13-1
Scroll bar show and hide, 8-3
Scrolling while sampling, 8-21
SD, 8-13
SDAC directive, 5-20
Search data, 11-3
Search for text, 7-7
 and replace it, 7-8
Selecting a channel, 2-1
Selecting areas in cursor windows, 11-6
Selecting channels, 2-1
SEM, 8-13
Send Mail, 6-10
Sequencer. *See* Output sequencer
Serial line input
 TextMark data, 4-8
Serial number, 14-2
SET directive, 5-20
Set Marker Codes, 9-41
Shell extensions, 3-6
Short cut keys
 Data views, 3-3
 Text views, 3-3
Show hidden window, 10-1
Show scroll bar, 8-3
Signal conditioner, 18-1
 Sample menu, 12-4
Sin() virtual channel function, 9-32
Sine wave output, **5-29**
Skip NaN, 9-40
Skyline display mode
 Result views, 8-13
 Time views, 8-12
Slope of line, 2-6, **11-7**
 trend plot, 9-18
Slope of wave, 9-39
Slope peak and trough, 11-5
Slope search cursor modes, 11-4
Slope% cursor mode, 11-5
Slow
 Drawing, break out of, 8-14
 Sampling configuration, 4-17
Smooth wave, 9-39
smr standard file extension, 6-1
Software help desk, 14-1
Son data storage library, 1-3
Son library, 1-3
SonFix, 4-25, 20-1
Sonogram colours, 8-20
Sonogram display mode, 8-11
sonview.tip file, 14-1
Sorted raster display, 8-14
Sound card, 12-3, 15-6
SpE() virtual channel function, 9-30
Spike Monitor, 8-24, 15-20
Spike sorting, **15-1**
 Create one channel per code, 15-14
 Detecting spikes, 15-1
 Digital output on match, 15-2
 Load templates, 15-18
 Off-line editing, 15-13
 Off-line sorting, 15-11
 On-line display, 15-17
 Sampling rate, 15-2
 Setup sampling for, 4-9
 Template controls, 15-6
 Template formation, 15-8
 Template setup, 15-3
Spike2 command line, 3-6
Spike2 data file format, 1-3
Spike2 for DOS scripts, 13-2
Spike2 Last configuration file, 4-30
Spike2 top box, 4-5
Spike2 versions, 1-3
Spikes per second
 Event correlation, 9-4
 PSTH, 9-3
Spreadsheet output, 7-2
Sqr() virtual channel function, 9-32
Sqrt() virtual channel function, 9-32
srf Result file extension, 6-1
Standard deviation
 between cursors, 11-8
 curve fitting, 9-25
 In result view, 8-13
 trend plot, 9-19
Standard display settings, 8-3
Standard error of the mean, 8-13
Start sampling, 4-28
startup.s2s, 3-6
State display mode, 8-13
Static cursor, 11-4
Status bar, 2-1
Stereotrode, 4-9
sTick() sequencer expression, 5-18
Stop band, 17-10, 17-11
Stop Process command, 9-12
Stop sampling, 4-28
SUB output instruction, 5-39
Sum of result view bins, **11-7**
Sweeps in a histogram, 8-4
sxy XY file extension, 6-1
System software required, 1-3
SZ output instruction, 5-30
SZINC output instruction, 5-30
- T**—
- TABADD output instruction, 5-41
TABDAT directive, 5-21
TABINC directive, 5-21
TABLD output instruction, 5-41
TabPos() sequencer expression, 5-19
TABST output instruction, 5-41
TABSUB output instruction, 5-41
TABSZ directive, 5-21
Talkers, 4-10
Tan() virtual channel function, 9-32
Template
 Copy to clipboard, 15-12
 Dialog, 15-9
 Digital output on match, 15-2
 Editing, 15-13
 Formation algorithm, 15-8
 Load, 15-18
 Merging, 15-5
 Off-line formation, 15-11
 On-line/Off-line, 15-17, 15-19
 Setting spike region, 15-4
 Setup window, 15-3
Ternary operator, 3-2
Tetrode, 4-9
Text display mode, 8-13
Text from different system, 7-1
Text output
 From result view, 7-2
 From time view, 7-3
 From XY view, 7-3
Text view control keys, 3-3
TextMark data, 4-8
 Copy As Text format, 7-6
 Copy to clipboard, 12-2
 Create during sampling, 12-1
 Export As format, 7-6
 Jump to marker in list, 12-2
 Serial input, 4-8
 Text display mode, 8-13
Threshold crossing cursor modes, 11-4
TICKS output instruction, 5-42
Tile windows, 10-1
Time at point, 9-18
Time difference, 9-18
Time expressions in dialogs, 3-2
Time in automatic file name, 4-24
Time limits, 4-24
Time resolution, 4-12
Time shift, 17-9
Time shift wave, 9-39
Time units per ADC convert, 4-12, 4-15
Time view, **2-1**, 2-8
 Copy as text, 7-2, 7-3

- Duplicate, 10-1
 - Time zero, 2-6, **11-6**
 - Timed sampling mode, 4-18
 - Timing faults in the graphical editor, 5-9
 - Tip of the Day, 14-1
 - Toolbar, 2-1
 - Play waveform, 4-20
 - Trend plot example, 9-15
 - Trigger channel
 - for correlations, 9-4, 9-5
 - for Phase histogram, 9-6
 - for PSTH, 9-3
 - for waveform average, 9-7
 - Trigger connections, 4-21
 - Rear panel, 7-16
 - Triggered displays, 8-6
 - Triggered sampling mode, 4-18
 - Keyboard marker trigger, 4-6
 - TextMark channel as trigger, 12-2
 - Triggered start of sampling, 4-28
 - Triggered waveform output, 4-21
 - Troubleshooting. *See* Try1401
 - Trough find cursor mode, 11-4
 - Try1401, 14-1
 - TTL compatible signals, 4-5
 - Turning point cursor mode, 11-5
 - Two band pass filter, 17-6
 - Two band stop filter, 17-6
 - txt text file extension, 6-1
- U—**
- Undo command, 7-1
 - Units for waveform data, 4-4
 - Un-magnify mouse pointer, 2-3
 - Updating Spike2, 1-4
 - us as time scaler, 3-2
 - us () sequencer expression, 5-18
 - usTick () sequencer expression, 5-18
 - Utility programs, 20-1
- V—**
- Valid cursors, 11-3
 - Value at cursor, 11-6
 - Value at point, 9-18
 - Value between cursors, 11-7
 - Value difference, 9-18
 - VAngle () sequencer expression, 5-19
 - VAR directive in output sequencer, 5-19
 - Variables for output sequencer, 5-19
 - VarValue script, 5-20, 5-24, 5-34
 - VDAC0-7 sequencer variables, 5-19
 - VDAC16 () sequencer expression, 5-19
 - VDAC32 () sequencer expression, 5-19
 - VDigIn sequencer variable, 5-19
 - Vertical Marker, 8-15
 - Vertical Markers
 - Line thickness, 7-15
 - VHz () sequencer expression, 5-19
 - Video codecs, 21-5
 - Video recording, 21-1
 - View handle
 - Interactive access to, 10-2
 - View menu, 8-1
 - Channel Draw Mode, 8-10
 - Colour commands, 8-20
 - Display trigger, 8-6
 - Enlarge and reduce, 8-1
 - File information for time view, 8-4
 - Font, 8-17
 - Info, 8-4
 - Rerun, 8-21
 - Result view drawing modes, 8-13
 - Show/Hide channel, 8-3
 - Spike Monitor, 8-24
 - Standard display, 8-3
 - X Axis Range, 8-1
 - Y Axis Range, 8-2
 - View-based expressions, 3-1
 - Virtual channels, 9-28
 - Voltage limits
 - 10 Volt/5 Volt ADC inputs, 7-17
 - TTL inputs, 4-5
 - VSz () sequencer expression, 5-19
- W—**
- WAIT output instruction, 5-26
 - Clash with template match, 15-2
 - WAITC output instruction, 5-33
 - WAVEBR output instruction, 5-45
 - Waveform data
 - 10 Volt/5 Volt range, 7-17
 - Aliasing, 4-4
 - Average, 9-7
 - Channel dialog, 4-3
 - Connections for sampling, 4-4
 - Convert to events, 9-28, 15-11
 - Convert to Marker, 15-11
 - Convert to WaveMark, 15-11
 - Copy As Text, Export As format, 7-4
 - Display mode, 8-12
 - Level crossing, 9-28
 - Peak search, 9-28
 - Power spectrum, 9-8
 - Sample rate, 4-4
 - Scaling, 4-4
 - Waveform output, 12-3
 - During sampling, 5-1
 - On-line, 4-20
 - start and stop from sequencer, 5-45
 - test from output sequencer, 5-45
 - WAVEGO output instruction, 5-44
 - WaveMark data, 15-1
 - Copy As Text, Export As format, 7-5
 - Detecting spikes, 15-1
 - Display mode, 8-12
 - From waveform, **15-11**
 - Overdraw display mode, 8-11
 - Reclassify, 15-13
 - Set codes in overdraw mode, 8-11
 - Set codes with line, 8-11, 8-13
 - Setup sampling for, 4-9
 - Spike sorting setup, 15-2
 - Traces, 9-26, 9-27
- WAVEST output instruction, 5-45
- Web site, 1-4
 - WEnv () virtual channel function, 9-31
 - Window for FFT, 9-9
 - Window menu, 10-1
 - Working Set, 14-2
 - WPoly () virtual channel function, 9-31
 - Write memory buffer to channel, 9-28
 - Write to disk, 4-18, 4-28
 - WSin () virtual channel function, 9-31
 - WSqu () virtual channel function, 9-31
 - WT () virtual channel function, 9-31
 - WTri () virtual channel function, 9-31
- X—**
- X axis control, 2-4
 - Short cut keys, 8-1, 8-2
 - Tick spacing, 8-1
 - Zero at trigger, 8-6
 - X Range dialog, 2-4, **8-1**
 - XOR, XORI output instructions, 5-40
 - XY Draw Mode, 8-14
 - XY view, 2-10
 - Auto-expand axes, 8-5
 - Copy as text, 7-3
 - Draw mode, 8-14
 - Example, 2-12
 - Fill with colour, 8-14
 - Key, 8-5
 - Line style, 8-14
 - Options, 8-5
 - Point style, 8-14
 - Prompt to save unsaved view, 7-11
- Y—**
- Y axis control, 2-5
 - No invert on drag, 7-16
 - Y Range dialog, 8-2
 - Y zero, 2-6, **11-6**
- Z—**
- Zero region, 2-6, **11-7**
 - Zero x axis at trigger, 8-6
 - Zoom in, out button, 2-2
 - Zoom text, 3-4