

# Dataguzzler as a Software Oscilloscope

Stephen D. Holland

September 4, 2007

Dataguzzler was written by Stephen D. Holland at the Iowa State University Center for Nondestructive Evaluation.

See the Dataguzzler web site at [http://ahab.cnde.iastate.edu/~sdh4/dg\\_web/](http://ahab.cnde.iastate.edu/~sdh4/dg_web/) for more information about Dataguzzler.

Dataguzzler is Copyright (C) 2005-2006 by Iowa State University.

Dataguzzler is released under the GPL 2.0/LGPL 2.1 licenses, with exceptions. Go to [http://ahab.cnde.iastate.edu/~sdh4/dg\\_web/dataguzzler/COPYING.txt](http://ahab.cnde.iastate.edu/~sdh4/dg_web/dataguzzler/COPYING.txt) for more information on TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION.

This material is based upon work supported by the Air Force Research Laboratory under Contract #FA8650-04-C-5228 at Iowa State University's Center for NDE.

This material is based upon work supported by the Federal Aviation Administration under Contract #DTFA03-98-D-00008, Delivery Order #0037 and performed at Iowa State University's Center for NDE as part of the Engine Titanium Consortium Phase III Thermal Acoustic Studies program.

Thanks to Ricky Reusser for designing the original logo and Kira Scott for drawing the new logo. Thanks to David Holland for suggestions on interfacing with the operating system.

Please take note of the warranty information as stated in the GPL v2.0 license:

#### NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

# Contents

<b>Introduction</b>	<b>6</b>
<b>1 Installation</b>	<b>7</b>
1.1 System requirements . . . . .	7
<b>2 System Control</b>	<b>8</b>
2.1 Manual control . . . . .	8
2.2 Oscilloscope display . . . . .	10
2.3 Utilities . . . . .	11
2.3.1 dg_save_settings . . . . .	11
2.3.2 dg_restore_settings . . . . .	11
2.3.3 dg_grab . . . . .	11
2.3.4 dg_grab.txt . . . . .	12
2.3.5 dg_cmd . . . . .	12
2.3.6 dg_upload . . . . .	12
2.3.7 dg_upload.txt . . . . .	12
2.3.8 dg_upload.tiff . . . . .	12

2.3.9	dg_snapshot . . . . .	13
2.3.10	dg_load_snapshot . . . . .	13
2.4	API Reference . . . . .	13
<b>3</b>	<b>System Operation</b>	<b>14</b>
<b>4</b>	<b>Command Reference</b>	<b>16</b>
	AUTH . . . . .	17
	DAC:GAIN . . . . .	18
	DAC:VOLTAGE . . . . .	19
	MATH:CLEARAVG . . . . .	20
	MATH:CLEARACCUM . . . . .	21
	MATH:DEF . . . . .	22
	MATH:ENABLE . . . . .	25
	MATH:ENABLED . . . . .	26
	MATH:DISABLE . . . . .	27
	MATH:UNDEF . . . . .	28
	MATH:UNDEFALL . . . . .	29
	MATH:WAITAVG . . . . .	30
	TIME:DELAY . . . . .	31
	TIME:TIMESTAMP . . . . .	32
	TRIG:MODE . . . . .	33
	TRIG:RATE . . . . .	34
	TRIG:TRIGGER . . . . .	35

WCAPT:ATRIGHIGH	36
WCAPT:ATRIGLOW	37
WCAPT:ATRIGMODE	38
WCAPT:CALCSYNC	39
WCAPT:CLKSRC	40
WCAPT:DSFACTOR	41
WCAPT:FIFOSIZE	42
WCAPT:FREQ	43
WCAPT:HWFREQ	44
WCAPT:HWTRIGSRC	45
WCAPT:NUMCHANNELS	46
WCAPT:CHi:PROBEATTEN	47
WCAPT:CHi:RANGE	48
WCAPT:SAMPLECNT	49
WFM:COPY	50
WFM:DATA	51
WFM:DATASHM	52
WFM:DELETEALL	53
WFM:DELETE	54
WFM:GLOBALREADYREV	55
WFM:GLOBALREADYREVTIMEOUT	56
WFM:GLOBALREV	57
WFM:GLOBALREVTIMEOUT	58

WFM:LIST	59
WFM:LISTREADY	60
WFM:LISTLOCK	61
WFM:LISTREADYLOCK	62
WFM:METADATA	63
WFM:REALSZ	64
WFM:REVISION	65
WFM:REVISIONLOCK	66
WFM:REVISIONREADYLOCK	67
WFM:UNLOCK	68
WFM:WFMS	69

**A Dataguzzler native binary file format 70**

A.1 Standardized file name extensions	72
A.2 File access API	72
A.2.1 Reading a Dataguzzler file	72
A.2.2 Writing a Dataguzzler file	74
A.2.3 MATLAB/Octave file access library	74
A.2.4 Python file access library	75

# Introduction

This manual describes the use of Dataguzzler as a software oscilloscope based on the PCI-DAS4020/12 waveform acquisition card. It is presented as an example of a simple, single-purpose Dataguzzler application. For more information on configuring Dataguzzler for more sophisticated applications, see the Dataguzzler manual: *Dataguzzler: A Generalized High Performance Data Acquisition Architecture*.

# Chapter 1

## Installation

### 1.1 System requirements

- Linux with 2.6 or newer kernel.
- The Measurement Computing PCI-DAS4020/12 waveform capture card is required.
- Warren Jasper's DAS4020 driver must be installed, with `pci-das4020.h` in `/usr/local/include`. You may want to increase the default buffer size `ADC_BUFF_PHY_SIZE` in `a2dc.h` before compiling the driver.

Dataguzzler should be compiled with the `make` command, and installed with the `make install` command. Watch the “make” output for errors.

In order to enable use of the Dataguzzler oscilloscope configuration by default, you may need to symbolic-link it to `dataguzzler.conf`:

```
rm -f /usr/local/dataguzzler/conf/dataguzzler.conf
ln -s dataguzzler_oscope.conf /usr/local/dataguzzler/conf/dataguzzler.conf
```



## Chapter 2

# System Control

The data acquisition system consists of a *dataguzzler server* containing a set of interconnected modules, plus various *clients* that connect to the server and issue commands and transfer data. Commands to the server can be issued from the dataguzzler console or over TCP/IP links. In contrast, the clients are separate programs that themselves connect to the server. To operate dataguzzler, first run the server in a terminal window (xterm):

```
pequod% dataguzzler
dataguzzler>
```

Commands to the server can be issued directly in that terminal window. Clients, on the other hand must be started separately. Open a new terminal window and type `dg_scope`. A oscilloscope display will appear. The oscilloscope display shows the current waveforms and images in dataguzzler memory. Note that it is an observation tool only; no changes or adjustments may be made from the oscilloscope display. Other clients are used to load and save settings and to upload and download waveforms from the server.

### 2.1 Manual control

Dataguzzler is command driven. Commands may be issued either on the terminal from which dataguzzler was run or over a TCP/IP connection. A TCP/IP connection can be established with the telnet command, e.g.:

```
linux% telnet localhost 1649
Trying 127.0.0.1...
Connected to marius (127.0.0.1).
Escape character is '^]'.
auth xyzyy
```

```
200 0000009 AUTH_OK
wcapt:freq?
200 0000019 WCAPT:FREQ 10 MHz
wcapt:freq 1 MHz
200 0000018 WCAPT:FREQ 1 MHz
quit
Connection closed by foreign host.
linux%
```

Let's examine the above transcript and learn how to issue commands. The text in **boldface** indicates what was typed by the user, while the **typewriter** text was generated by the computer.

Upon connecting to port 1649 (the dataguzzler port), the first command issued was **auth xyzzy**. This authenticates the user to the dataguzzler server and must be done before any commands can be issued (authentication is not necessary when typing commands on the console). Information on the **AUTH** command can be found in chapter 4 on page 17.

Dataguzzler then responds initially with exactly 17 characters. The first 3 characters are the return code in decimal. 200 indicates success, 500 or higher indicates an error. The return code is followed by a space and 12 more characters which indicate the length of the response (not including the 17 character header), again in decimal. This is followed by a space, then the remainder of the response, with the length as specified. The response ends in a carriage return and linefeed. These characters are included in the length count. In the case of the **AUTH** command, the specified length was 9 characters. The actual response was "AUTH\_OK", 7 characters, plus the carriage return and linefeed, for a total of 9.

The AUTH\_OK response indicates that authentication has been successfully completed and that other commands may be issued. The user then queries the waveform capture capture sample frequency (WCAPT:FREQ, pg. 43), adjusts the capture frequency from 10 MHz to 1 Mhz, and exits.

The complete list of available commands can be found in chapter 4. It is important to realize that the oscilloscope display program has no control capabilities whatsoever. Its sole purpose is to display the waveforms in the dataguzzler memory. All acquisition parameter changes must be performed separately.

Normally each command transmitted is terminated by a carriage return, linefeed, or combination thereof. It is possible to transmit a number of commands as a single atomic unit by separating them with semicolons instead. The composite command must still have a linebreak at the end. The reply generated by the server will consist of a single block containing the corresponding responses similarly separated by semicolons and with a carriage return / linefeed pair at the end. In this situation all the specified commands will be executed together as an atomic unit unless waiting is required by one of the commands. Commands that wait are documented as such in chapter 4. If an error occurs while executing a command, the command's reply may be replaced by an error message. The substring "ERROR" should occur in this message before the first space.

## 2.2 Oscilloscope display

The oscilloscope display can be started with the `dg_scope` command. The oscilloscope display assumes dataguzzler is running on host `localhost` port `1649` with authentication code `xyzy`. If dataguzzler is running elsewhere, these parameters can be provided on the command line:

```
scope <hostname> <port> <authcode>
```

Additional parameters defined by X and GLUT can also be provided. See <http://www.opengl.org/developers/documentation/glut/spec3/node10.html> for a complete list

The oscilloscope display provides live viewing of the waveforms in the dataguzzler memory, and real-time manipulation of that view. *The oscilloscope display is a tool only for viewing the dataguzzler waveforms and settings. The oscilloscope display cannot be used to change the dataguzzler settings.*

The oscilloscope display is designed for keyboard or combined mouse and keyboard interaction. The keyboard commands are as follows:

Enter	Disable or enable display selected waveform
Tab	Select next waveform
Cursor left	Increase Secs/Div, Hz/Div, or pixels/pixel (zoom out horizontally)
Cursor right	Decrease Secs/Div, Hz/Div, or pixels/pixel (zoom in horizontally)
Cursor down	Increase Volts/Div. (zoom out vertically) or decrease image contrast.
Cursor up	Decrease Volts/Div (zoom in vertically) or increase image contrast.
Home	Increase $t_0$ by one division (look later in waveform)
End	Decrease $t_0$ by one division (look earlier in waveform)
PgUp	Increase display offset of waveform by one vertical division
PgDn	Decrease display offset of waveform by one vertical division
Insert	Increase image brightness
Delete	Decrease image brightness
'>	Select previous frame of a multi-frame image.
'<	Select next frame of a multi-frame image.
'c'	Cycle between colormaps for image displays
'o'	Set the position or offset to the median value of the selected waveform or frame
'z'	Zero the position and offset of the selected waveform or frame

The oscilloscope window consists of the active oscilloscope area, surrounded by informational displays. The left edge of the window lists the names of the waveforms available from the server. The currently selected waveform name is highlighted. Clicking on a waveform name will select that waveform. If an attenuating probe was used and configured in the server that information (e.g. 10x) will be recorded after the waveform name. Next to the waveform name is a clickable box which controls whether that waveform is currently displayed.

The top line of the window provides the acquisition parameters for the currently selected waveform.  $t_0$  is the time of the first sample. Fs is the sample frequency. n is the total number of samples acquired, and rev is the waveform revision. For AVG and AVGNCE waveforms, the status of the averaging is also indicated.

The bottom of the display lists the current display parameters.  $t_0$  is the time corresponding to the vertical line in the center of the waveform display.  $f_0$  is the frequency corresponding to the vertical line in the center of the waveform display (for frequency domain waveforms). `globalrev` is the current overall waveform revision count. `Position` is the vertical offset of the selected waveform.

The waveform display uses several techniques to reduce aliasing artifacts and improve signal comprehension. If there are more than two waveform points corresponding to a particular vertical line of pixels on the display, then a vertical line is drawn between the extreme pixels. For less than two points, the samples themselves are drawn. If there is no waveform point corresponding to a particular point on the display, an interpolated value is drawn in a lighter color. The user should be warned that the interpolation algorithm is approximate only. Proper interpolation can be performed in post-processing using algorithms such as Matlab `INTFILT`.

A standalone version of the oscilloscope display can be run with the command `dg_scope_sa`. This version does not require `dataguzzler` to be running, but instead displays one or more waveform file (`.dgz`) or snapshot file (`.dgs`) specified on the command line.

## 2.3 Utilities

### 2.3.1 `dg_save_settings`

`dg_save_settings` is an external program that downloads the current state of `dataguzzler` and writes it to a file.

The general usage is:

```
dg_save_settings <settings_file.set>
```

Additionally `-h <host_name>` and `-a <authentication_code>` parameters may be provided.

`dg_save_settings` stores the results of the `WFM:WFMS?` and `SET?` commands to the output file.

### 2.3.2 `dg_restore_settings`

`dg_restore_settings` is an external program that uploads a stored settings file to `dataguzzler`. The general usage is:

```
dg_restore_settings <settings_file.set>
```

Additionally `-h <host_name>` and `-a <authentication_code>` parameters may be provided.

`dg_restore_settings` assumes the file consists of two lines of commands. It passes both lines to `dataguzzler` to be executed.

### 2.3.3 `dg_grab`

`dg_grab` is an external program that downloads waveforms from `dataguzzler` and writes them to `dataguzzler` format binary files. See Appendix A for more information on the file format. The general usage is:

`dg_grab` <waveform\_name> <file\_name> <waveform\_name2> <file\_name2> ...  
Additionally `-h` <host\_name> and `-a` <authentication\_code> parameters may be provided.  
As many (waveform\_name, file\_name) pairs as are desired can be specified.

### 2.3.4 dg\_grab\_txt

`dg_grab_txt` is an external program that downloads waveforms from dataguzzler and writes them to ASCII text files. The general usage is:

`dg_grab_txt` <waveform\_name> <file\_name> <waveform\_name2> <file\_name2> ...  
Additionally `-h` <host\_name> and `-a` <authentication\_code> parameters may be provided.  
As many (waveform\_name, file\_name) pairs as are desired can be specified.

### 2.3.5 dg\_cmd

`dg_cmd` issues a single command to dataguzzler and prints the response to stdout. The general usage is:

`dg_cmd` <command>  
Additionally `-h` <host\_name> and `-a` <authentication\_code> parameters may be provided.

### 2.3.6 dg\_upload

`dg_upload` uploads a binary dataguzzler format file to dataguzzler. See Appendix A for more information on the file format. The general usage is:

`dg_upload` <waveform\_name> <file\_name> <waveform\_name2> <file\_name2> ...  
Additionally `-h` <host\_name> and `-a` <authentication\_code> parameters may be provided.  
As many (waveform\_name, file\_name) pairs as are desired can be specified.

### 2.3.7 dg\_upload\_txt

`dg_upload_txt` uploads an ASCII text waveform to dataguzzler. The general usage is:

`dg_upload_txt` <waveform\_name> <file\_name> <waveform\_name2> <file\_name2> ...  
Additionally `-h` <host\_name> and `-a` <authentication\_code> parameters may be provided.  
As many (waveform\_name, file\_name) pairs as are desired can be specified.

### 2.3.8 dg\_upload\_tiff

`dg_upload_tiff` uploads a TIFF image to dataguzzler as a grayscale. The general usage is:

`dg_upload_tiff` <waveform\_name> <file\_name> <waveform\_name2> <file\_name2> ...

Additionally `-h` <host\_name> and `-a` <authentication\_code> parameters may be provided.

As many (waveform\_name, file\_name) pairs as are desired can be specified. Note: `dg_upload_tiff` cannot read grayscale images with more than 8 bits per pixel.

### 2.3.9 dg\_snapshot

`dg_snapshot` saves a consistent snapshot of all waveforms (including dynamic waveforms) to the name (.dgs file) specified on the command line.

### 2.3.10 dg\_load\_snapshot

`dg_load_snapshot` loads the waveforms in the .dgs file specified on the command line into dataguzzler. Please note the potential for conflicts with pre-existing channels. Depending on the nature of the pre-existing channel, the conflict may be resolved one way or the other. In general, only nonexistant channels or preexisting channels owned by the WFMIO module can be overwritten by `dg_load_snapshot`.

## 2.4 API Reference

A library of C utility functions is installed in `/usr/local/dataguzzler/lib`, with include files in `/usr/local/include`. A similar library of functions for Matlab or GNU Octave is installed in `/usr/local/dataguzzler/matlab`. These libraries are used for convenient access to the dataguzzler server and provide routines for remote access and waveform upload/download. Both APIs lack formal documentation at this time. Nevertheless, as the APIs are simple and straightforward it should not be difficult to learn them anyway. All functions are prototyped in the include files, and the source code for the utilities described above make excellent examples for the C library. For the Matlab/Octave routines see the files `dgf_testread.m` and `dgf_testwrite.m` for simple examples, and `proccalib.m` for a more complicated example.

## Chapter 3

# System Operation

The Dataguzzler-based oscilloscope, because of its simplicity, has only a few important parameters. `WCAPT:SAMPLECNT` is used to set the number of samples to acquire per trigger. Please note that the maximum `samplecnt` is limited by the size of the DAS4020 driver's buffer. This can be changed by adjusting `ADC.BUFF.PHY.SIZE` in `a2dc.h` and recompiling the DAS4020 driver.

`WCAPT:FREQ` is used to set the sample rate and `WCAPT:NUMCHANNELS` sets the number of acquisition channels. `WCAPT:CH1:RANGE` through `WCAPT:CH4:RANGE` set the different channels to the 1V or 5V input ranges.

By default, the oscilloscope is triggered from Channel 1 with hysteresis such that the voltage on channel 1 must drop *below* .2 Volts, then must *exceed* .7 volts in order to generate a trigger. The thresholds may be adjusted with `WCAPT:ATRIGLOW` and `WCAPT:ATRIGHIGH`.

The trigger source may be selected with `WCAPT:HWTRIGSRC`:

- `INT`: pin 10 (PC5) of the IDC header on the DAS4020 card.
- `EXT`: bottom BNC connector on the card, and
- `CH1` through `CH4`: analog triggering on each of the four input channels.

A software trigger generator module is included. It generates its triggers on pin PB5 of the IDC header. The default frequency is 10 Hz. Therefore, if computer-generated or "auto" triggering is required, you can jumper together pins 9 and 10 (PB5 and PC5) of the IDC header and set `WCAPT:HWTRIGSRC` to `INT`. Then you will be able to adjust the auto-trigger repetition rate with `TRIG:RATE`, switch to or from manual trigger generation with `TRIG:MODE`, and manually generate triggers with `TRIG:TRIGGER`.

The digital to analog converters on the DAS4020 are supported with the DAC module and the `DAC:VOLTAGE` command.

Waveform transfers are supported by the WFM module and realtime computation is supported by the MATH module.



## Chapter 4

# Command Reference

A wealth of commands are provided to communicate with the hardware devices. Each command begins with the name of the corresponding module. These commands can be given either from the console or over a TCP/IP connection. In addition to these commands, `SET?` queries the status of all settings, and `SET` causes all modules to rewrite settings to hardware.

(commands begin on next page)

# AUTH

## Syntax

AUTH <authcode>

## Description

Authenticates an incoming TCP/IP connection

## Parameters

<authcode>                      Authentication password

## Notes

- You must use an authentication code that is valid for the IP address you are connecting from. Authentication is configured using the file `/etc/daq_auth.conf`. If this file does not exist, internal defaults allow connections only on the loopback address (127.0.0.1) using `xyzy` as the authentication code.

## See also

# DAC:GAIN

## Syntax

```
DAC:GAIN<channel number> <gain setting>  
GAIN<channel number>?
```

## Description

Specify or query the gain setting of DAC channel <channel number>.

## Parameters

<channel number>	The channel number for which to set or query the gain.
<gain setting>	The gain setting, 1V or 10V

## Notes

- When the gain is changed, the output voltage on that channel will briefly drop to 0V.
- Capacitive loading such as a long cable may cause the DAC output amplifier to become unstable for the 1V gain setting.
- The default gain is 10V.
- <channel number> may be 0 or 1.

## See also

- DAC:VOLTAGE (pg. 19)

# DAC:VOLTAGE

## Syntax

```
DAC:VOLTAGE<channel number> <voltage>  
VOLTAGE<channel number>?
```

## Description

Specify or query the output voltage of DAC channel <channel number>.

## Parameters

<channel number>	The channel number for which to set or query the voltage.
<voltage>	The voltage to set (default units of Volts)

## Notes

- <channel number> may be 0 or 1

## See also

- DAC:GAIN (pg. 18)

# MATH: CLEARAVG

## Syntax

MATH: CLEARAVG <waveform name>

## Description

Reset averaging channel <waveform name>

## Parameters

<waveform name>      Name of the math waveform to reset.

## Notes

## See also

# MATH: CLEARACCUM

## Syntax

MATH: CLEARACCUM <waveform name>

## Description

Reset ACCUM channel <waveform name>

## Parameters

<waveform name>      Name of the MATH:ACCUM waveform to reset.

## Notes

- The the ACCUM channel will be reset and will be empty until a new version of the waveform it is dependent on appears.

## See also

# MATH:DEF

## Syntax

MATH:DEF <waveform name>=<function name>(<parameters ... >)  
DEF? <waveform name>

## Description

Define a new channel to be a mathematical function of other channel(s), or query the mathematical function of such a channel

## Parameters

<waveform name>	Name of the waveform to define or query.
<function name>	Name of mathematical function to use
<parameters ... >	Parameters to the mathematical function

## Notes

- Allowable functions are:
  - result=AVG(<channel name>,<number of averages>) or
  - (result,stddev)=AVG(<channel name>,<number of averages>): Running average (and optional standard deviation) of <channel name>.
  - result=AVGONCE(<channel name>,<number of averages>) or
  - (result,stddev)=AVGONCE(<channel name>,<number of averages>): Average of (and optional standard deviation) of <channel name>.
  - ampl=FFT(<channel name>,<transform dimensions>) or
  - (ampl,phase)=FFT(<channel name>,<transform dimensions>): Fourier Transform of <channel name>. If <transform dimensions> is not specified the transform will be over the first (minor) dimension of the data. <transform dimensions> can either be an integer, specifying which dimension to transform over (0 being the first – minor – dimension), or it can be a list of comma separated such integers enclosed within square brackets, in which case the transform will be over all the specified dimensions, e.g. FFT(chan1,[0,2,3]) will cause a Fourier transform over the first, third, and fourth dimension. The highest transformed dimension (fourth dimension in the previous example) will have size  $(n/2) + 1$  where n is the pre-existing length of that dimensions and the result of the division is truncated to the next lower integer, not rounded. All other dimensions will have their pre-existing sizes. There can be either one or two result parameters. The first (or only) result parameter is the amplitude of the Fourier transform. The second result parameter is the phase (in radians) of the

Fourier transform. The transform is normalized by multiplying by the product of the step sizes of the transformed dimensions. .

- result=CORR(<channel 1>,<channel 2>,<dimensions>), or
- result=CONV(<channel 1>,<channel 2>,<dimensions>): CORR and CONV perform cross-correlation and convolution respectively of <channel 1> with <channel 2> over dimensions <dimensions>. <dimensions> can be an integer, specifying which dimension to correlate/convolve over (0 being the first – minor – dimension), or it can be a list of comma separated such integers enclosed within square brackets, in which case the correlation/convolution will be over all the specified dimensions. If <dimensions> is not specified, it will be over the first (minor) dimension.
- result=ACCUM(<channel name>,<number of waveforms>): Accumulate a series of waveforms into a “waveform” with an extra dimension. For example, if <channel name> is two-dimensional 640x512 and <number of waveforms> is 6 then the result will be a “cube” of data 640\*512\*6. Note that all the input waveforms must be the same size or the generated result will be empty. Also note that ACCUM will not include the current revision in its series, but will start accumulating with the next version.
- result=ACCUMONCE(<channel name>,<number of waveforms>): Like ACCUM but doesn’t automatically reset when full. Need to call MATH:CLEARACCUM manually.
- result=ADD(<channel a>,<real number b>) or
- result=ADD(<channel a>,<channel b>): Compute result=a+b. b may be of lower dimensionality than a, but the dimensions of b must match the first dimensions of a. The result gets a copy of a’s metadata.
- result=SUB(<channel a>,<real number b>) or
- result=SUB(<channel a>,<channel b>): Compute result=a-b. b may be of lower dimensionality than a, but the dimensions of b must match the first dimensions of a. The result gets a copy of a’s metadata.
- result=MUL(<channel a>,<real number b>) or
- result=MUL(<channel a>,<channel b>): Compute result=a\*b. b may be of lower dimensionality than a, but the dimensions of b must match the first dimensions of a. The result gets a copy of a’s metadata.
- result=DIV(<channel a>,<real number b>) or
- result=DIV(<channel a>,<channel b>): Compute result=a/b. b may be of lower dimensionality than a, but the dimensions of b must match the first dimensions of a. The result gets a copy of a’s metadata.
- result=INT(<channel>) or
- result=INT(<channel>,<dimension>): Integrate <channel> over <dimension> (default 1). The result has the same dimensionality as <channel> and is the running integral over the specified dimension.
- result=DIFF(<channel>) or
- result=DIFF(<channel>,<dimension>): Differentiate <channel> over <dimension> (default 1). The result has the same dimensionality as <channel> and is the derivative over the specified dimension. The resulting waveform will be shifted by exactly 1/2 sample in the forward direction and its last sample will have a value of 0.0.



- result=INTEGRAL(<channel>) or
- result=INTEGRAL(<channel>,<dimension>): Calculate the integral of <channel> over the specified dimension (default 1). The result will have one less dimension than <channel> and is the integral over all the samples along the specified dimension.
- result=SUBEARLYAVG(<channel>,<threshold>) or
- result=SUBEARLYAVG(<channel>,<threshold>,<dimension>): Subtract the average of the first elements (up to <threshold>) along <dimension> of <channel> from <channel>. <dimension> is from 0 (minor dimension) to (ndim-1) (major dimension), with the major dimension used by default.
- result=FILTEREDINT(<channel>,<freq 1>,<freq 2>): Integrate the one-dimensional waveform in <channel>, then subtract out its average slope, then high-pass filter it with a frequency-domain raised-cosine that starts at 0.0 at <freq 1> and reaches 1.0 at <freq 2>.
- result=MAX(<channel>): Find the scalar maximum value (over all dimensions) of the specified channel.
- result=CROP(<channel>,[axis1min,axis1max],[axis2min,axis2max],...): Crop the specified waveform or image.
- result=DECIMATE(<channel>,<first axis decimate factor>,<second axis decimate factor>, ...): Downsample the specified channel by the specified factors in each axis.
- result=DBABS(<channel>): Convert to dB (return  $20 \log_{10}(\text{abs}(\text{<channel>}))$ )

## See also

- MATH:UNDEF (pg. 28)

# MATH:ENABLE

## Syntax

MATH:ENABLE <waveform name>

## Description

Enable math channel <waveform name>

## Parameters

<waveform name>      Name of the math channel to enable.

## Notes

- Math channels are enabled by default.
- Enabling one result channel of a function with multiple outputs enables all result channels of that function

## See also

- MATH:ENABLED (pg. 26)
- MATH:DISABLE (pg. 27)

# MATH:ENABLED

## Syntax

MATH:ENABLED? <waveform name>

## Description

Determine whether math channel <waveform name> is enabled.

## Parameters

<waveform name>      Name of the math channel to check.

## Notes

- Returns response of the form MATH:ENABLE <waveform name> or MATH:DISABLE <waveform name> depending on whether <waveform name> is enabled.

## See also

- MATH:ENABLE (pg. 25)
- MATH:DISABLE (pg. 27)

# MATH:DISABLE

## Syntax

MATH:DISABLE <waveform name>

## Description

Disable math channel <waveform name>

## Parameters

<waveform name>      Name of the math channel to disable.

## Notes

- Math channels are enabled by default.
- Disabling one result channel of a function with multiple outputs disables all result channels of that function

## See also

- MATH:ENABLE (pg. 25)
- MATH:ENABLED (pg. 26)

# MATH:UNDEF

## Syntax

MATH:UNDEF <math channel name>

## Description

Remove the math channel <math channel name>

## Parameters

<math name>	channel	Name of the math channel to remove.
----------------	---------	-------------------------------------

## Notes

## See also

- MATH:DEF (pg. 22)
- MATH:UNDEFALL (pg. 29)

# MATH:UNDEFALL

## Syntax

MATH:UNDEFALL <math channel name>

## Description

Delete all math channels

## Parameters

## Notes

## See also

- MATH:DEF (pg. 22)
- MATH:UNDEF (pg. 28)

# MATH:WAITAVG

## Syntax

MATH:WAITAVG <waveform name>

## Description

Wait for averaging channel <waveform name> to have performed a complete set of averages

## Parameters

<waveform name>      Name of the math averaging channel to wait for.

## Notes

- This waits for a *complete* and *ready* averaging channel. It does not lock the completed version in memory.

## See also

# TIME:DELAY

## Syntax

```
TIME:DELAY <time>  
DELAY?
```

## Description

Wait for a specified amount of time.

## Parameters

<time>                      The amount of time to wait, in units of time (default seconds)

## Notes

- This command does not occupy the dataguzzler kernel. Commands issued from other connections can execute during the wait.

## See also



# TIME:TIMESTAMP

## Syntax

TIME:TIMESTAMP?

## Description

Obtain a timestamp

## Parameters

## Notes

- The form of the result is a single quoted string, e.g. "2007-06-27T18:43:59-0500". This is believed to be ISO-8601 compliant.

## See also

# TRIG:MODE

## Syntax

```
TRIG:MODE <trigger mode>  
MODE?
```

## Description

Set the mode for the trigger generator.

## Parameters

<trigger mode>            The desired trigger mode: INTernal or COMPuter

## Notes

## See also

- TRIG:RATE (pg. 34)

# TRIG:RATE

## Syntax

TRIG:RATE <trigger frequency>

## Description

Set the frequency of the trigger generator.

## Parameters

<trigger frequency>      The desired trigger frequency (default units of Hz)

## Notes

- This frequency controls the minimum period between triggers (TRIG:MODE COMPUTER), or acts as an upper bound on the frequency of internally generated triggers (TRIG:MODE INTERNAL)/

## See also

- TRIG:TRIGGER (pg. 35)
- TRIG:MODE (pg. 33)

# TRIG:TRIGGER

## Syntax

TRIG:TRIGGER

## Description

Issue a trigger on the configured digital output.

## Parameters

## Notes

- Triggers will only be generated as fast as specified in TRIG:RATE. If you attempt to retrigger too quickly the TRIG:TRIGGER command will wait until the trigger period has passed before generating the trigger and returning.
- The trigger is generated on port B5 on the IDC header of the DAS4020 card. Place a jumper between port B5 and port C5 and select WCAPT:HWTRIGSRC INT to use this command to manually trigger the waveform capture card.

## See also

- TRIG:RATE (pg. 34)
- WCAPT:HWTRIGSRC (pg. 45)

# WCAPT:ATRIGHIGH

## Syntax

```
WCAPT:ATRIGHIGH <trigger voltage>  
ATRIGHIGH? <trigger voltage>
```

## Description

Specify or query the low analog trigger voltage.

## Parameters

<trigger voltage>      Desired trigger voltage

## Notes

- The high analog trigger voltage is used for rising-edge triggering, the hysteresis of falling edge triggering, and window triggering.
- Changing the analog trigger voltage will cancel any acquisition currently in progress on the capture card.
- The actual quantized trigger voltage will be returned
- Changing the gain setting or trigger source may cause the the trigger level to be re-quantized to match the new setting.

## See also

- WCAPT:HWTRIGSRC (pg. 45)
- WCAPT:ATRIGLOW (pg. 37)
- WCAPT:ATRIGMODE (pg. 38)

# WCAPT:ATRIGLOW

## Syntax

```
WCAPT:ATRIGLOW <trigger voltage>  
ATRIGLOW? <trigger voltage>
```

## Description

Specify or query the low analog trigger voltage.

## Parameters

<trigger voltage>      Desired trigger voltage

## Notes

- The low analog trigger voltage is used for falling-edge triggering, the hysteresis of rising edge triggering, and window triggering.
- Changing the analog trigger voltage will cancel any acquisition currently in progress on the capture card.
- The actual quantized trigger voltage will be returned
- Changing the gain setting or trigger source may cause the the trigger level to be re-quantized to match the new setting.

## See also

- WCAPT:HWTRIGSRC (pg. 45)
- WCAPT:ATRIGHIGH (pg. 36)
- WCAPT:ATRIGMODE (pg. 38)

# WCAPT:ATRIGMODE

## Syntax

```
WCAPT:ATRIGMODE <trigger mode>  
ATRIGMODE? <trigger mode>
```

## Description

Specify or query the analog trigger slope/mode.

## Parameters

<trigger mode>            Desired trigger mode

## Notes

Valid trigger modes are:

- **POS\_HIST** Trigger on analog voltage rising above ATRIGHIGH voltage with ATRIGLOW hysteresis (i.e. Schmitt trigger)
- **NEG\_HIST** Trigger on analog voltage falling below ATRIGLOW voltage with ATRIGHIGH hysteresis (i.e. Schmitt trigger)
- **POS\_SLOPE** Trigger on analog voltage rising above ATRIGHIGH voltage.
- **NEG\_SLOPE** Trigger on analog voltage falling below ATRIGLOW voltage.
- **WINDOW** Trigger on analog voltage between ATRIGLOW and ATRIGHIGH voltages.

Changing the analog trigger mode will cancel any acquisition currently in progress on the capture card.

## See also

- • WCAPT:HWTRIGSRC (pg. 45)
- WCAPT:ATRIGHIGH (pg. 36)
- WCAPT:ATRIGLOW (pg. 37)

# WCAPT:CALCSYNC

## Syntax

```
WCAPT:CALCSYNC <sync_enabled>  
CALCSYNC? <sync_enabled>
```

## Description

Set whether new acquisitions should be inhibited until computation from the previous acquisition is complete

## Parameters

<sync\_enabled> Whether new acquisitions should be inhibited, **true** or **false**

## Notes

- Effective following the next trigger. Computations currently in progress will not be waited for before allowing a trigger.

## See also



# WCAPT:CLKSRC

## Syntax

```
WCAPT:CLKSRC <clock source> Hz  
CLKSRC?
```

## Description

Specify or query the clock source for waveform capture A/D

## Parameters

<clock source>            Desired clock source

## Notes

- <clock source> may be INTERNAL, EXTBNC, or ADSTARTTRIG to select the internal 40 MHz source, the bottom BNC connector, or the A/D Start Trig pin on the IDC header, respectively.
- Changing the clock source will cancel any acquisition currently in progress on the capture card.
- The card is designed for a maximum clock rate of 40 MHz.
- Internal limitations of the card require divisors of at least 2 (1 or 2 channels) or 4 (4 channels), so depending on the number of channels the maximum sample rate will be 1/2 or 1/4 of the frequency of the external clock.
- Use WCAPT:HWFREQ to set the divisor and WCAPT:FREQ to set the actual divided clock frequency when not in INTERNAL mode.

## See also

- WCAPT:SAMPLECNT (pg. 49)
- WCAPT:NUMCHANNELS (pg. 46)
- WCAPT:FREQ (pg. 43)
- WCAPT:HWFREQ (pg. 44)

# WCAPT:DSFACTOR

## Syntax

```
WCAPT:CH<i>:DSFACTOR <factor>  
CH<i>:DSFACTOR?
```

## Description

Specify or query the downsampling factor for channel <i>

## Parameters

<factor>                      Desired downsampling factor.

## Notes

- If the downsampling factor is  $d$ , only one of every  $d$  samples will be recorded, and the time step will be  $d$  times the actual sample period.
- This is useful primarily for waveforms that are slow changing, or if due to unavoidable clocking restraints it is necessary to oversample.

## See also

- WCAPT:SAMPLECNT (pg. 49)
- WCAPT:NUMCHANNELS (pg. 46)
- WCAPT:FREQ (pg. 43)
- WCAPT:HWFREQ (pg. 44)

# WCAPT:FIFOSIZE

## Syntax

```
WCAPT:FIFOSIZE <fifo size>  
FIFOSIZE? <fifo size>
```

## Description

Specify or query the amount of the DAS internal FIFO to use.

## Parameters

<fifo size>                      Desired FIFO size, in 24-bit words

## Notes

- The default value of 32768 corresponds to 32768 24 bit words in each of the two (X and Y) FIFOs, for a memory-equivalent of 256 kilobytes.
- This may be set lower to avoid conflicts with other devices that have smaller FIFO's. By reducing the FIFO size, data transmissions are sent in smaller chunks that are less likely to overflow the FIFO of another device that is writing data to memory at the same time.
- Only powers of 2 starting at 256 and going up to 32768 are permitted

## See also

- WCAPT:SAMPLECNT (pg. 49)
- WCAPT:NUMCHANNELS (pg. 46)

# WCAPT:FREQ

## Syntax

```
WCAPT:FREQ <sample frequency> Hz  
FREQ? <sample frequency> Hz
```

## Description

Specify or query the number of samples to acquire per second

## Parameters

<sample frequency>	Desired sample rate (default Hz)
-----------------------	----------------------------------

## Notes

- Changing the sample frequency will cancel any acquisition currently in progress on the capture card.
- Only integer fractions of 20 MHz are permitted. When specifying the sample frequency, check the response from the server to determine the actual sample rate.
- The maximum sample rate is 20 MHz (1 or 2 channels) or 10 MHz (4 channels)
- If CLKSRC is not INTERNAL then this must be set to the actual divided clock frequency. In that case the clock division is determined by WCAPT:HWFREQ.

## See also

- WCAPT:SAMPLECNT (pg. 49)
- WCAPT:NUMCHANNELS (pg. 46)
- WCAPT:HWFREQ (pg. 44)

# WCAPT:HWFREQ

## Syntax

```
WCAPT:HWFREQ <sample frequency> Hz  
HWFREQ? <sample frequency> Hz
```

## Description

Specify or query the capture frequency to program the DAS4020 card with

## Parameters

<sample frequency>                      Desired capture frequency

## Notes

- Changing the sample frequency will cancel any acquisition currently in progress on the capture card.
- Only integer fractions of 20 MHz are permitted. When specifying the sample frequency, check the response from the server to determine the actual sample rate.
- The maximum HWFREQ is 20 MHz (1 or 2 channels) or 10 MHz (4 channels)
- If CLKSRC is INTERNAL, this will not be adjustable and will track WCAPT:FREQ.
- If CLKSRC is not INTERNAL then this is used to determine the clock division of the external clock. Program this with the frequency you would use to get the desired division if the external clock were 40 MHz. For example, to get a division of 4, program this to 10 MHz. To get a division of 12, program this to 3.33 MHz.

## See also

- WCAPT:SAMPLECNT (pg. 49)
- WCAPT:NUMCHANNELS (pg. 46)
- WCAPT:HWFREQ (pg. 44)

# WCAPT:HWTRIGSRC

## Syntax

```
WCAPT:HWTRIGSRC <trigger source>  
HWTRIGSRC? <trigger source>
```

## Description

Specify or query the trigger source

## Parameters

<trigger source>	Desired trigger source: “EXT” for the bottom BNC connector on the card or “INT” for the trigger line on the 40 pin header, or “CH1” through “CH4” for triggering from the analog input channels.
------------------	--

## Notes

- Changing the trigger source will cancel any acquisition currently in progress on the capture card.

## See also

- WCAPT:ATRIGMODE (pg. 38)
- TRIG:MODE (pg. 33)

# WCAPT:NUMCHANNELS

## Syntax

```
WCAPT:NUMCHANNELS <number of channels>  
NUMCHANNELS? <number of channels>
```

## Description

Specify or query the number of channels to acquire

## Parameters

<number of channels> of Desired number of channels

## Notes

- Changing the number of channels will cancel any acquisition currently in progress on the capture card.
- 1, 2, or 4 channels are allowed.
- A sample rate of 20MHz requires 2 or fewer channels

## See also

- WCAPT:SAMPLECNT (pg. 49)
- WCAPT:FREQ (pg. 43)

# WCAPT:CH*i*:PROBEATTEN

## Syntax

```
WCAPT:CH<i>:PROBEATTEN <attenuation factor>  
CH<i>:PROBEATTEN? <attenuation factor>
```

## Description

Specify or query the attenuation factor of the probe attached to channel <i>.

## Parameters

<i>	Channel number: 1-4
<attenuation factor>	Desired attenuation factor (default 1.0)

## Notes

## See also

- WCAPT:CH*i*:RANGE (pg. 48)



# WCAPT:CHi:RANGE

## Syntax

```
WCAPT:CH<i>:RANGE <input range>  
CH<i>:RANGE? <input range>
```

## Description

Specify or query input gain/attenuation setting for the capture card

## Parameters

<i>	Channel number: 1-4
<input range>	Desired input range: 1V or 5V

## Notes

- Changing the input range will cancel any acquisition currently in progress on the capture card.
- Current (as of Jan 2006) versions of the driver do not correctly handle the case of different gain settings on different channels. Hopefully this will be fixed eventually. In the mean time, be sure to use the same setting on all four channels.

## See also

- WCAPT:CHi:PROBEATTEN (pg. 47)

# WCAPT:SAMPLECNT

## Syntax

```
WCAPT:SAMPLECNT <number of samples>  
SAMPLECNT? <number of samples>
```

## Description

Specify or query the number of samples to acquire per waveform

## Parameters

<number of samples>      of      Number of samples to acquire

## Notes

- Changing the number of samples will cancel any acquisition currently in progress on the capture card.
- The maximum permissible SAMPLECNT is determined when the das4020 driver is compiled by the symbol ADC\_BUFF\_PHY\_SIZE specified in a2dc.h. To increase ADC\_BUFF\_PHY\_SIZE you must change that parameter and recompile and reinstall the driver.

## See also

- WCAPT:FREQ (pg. 43)

# WFM: COPY

## Syntax

WFM: COPY **<waveform name>** **<copy name>**

## Description

Copy the specified waveform.

## Parameters

<b>&lt;waveform name&gt;</b>	Name of the original waveform
<b>&lt;copy name&gt;</b>	Name for the copy

## Notes

- The copy can be deleted with the WFM:DELETE command.

## See also

- WFM:DELETE (pg. 54)

# WFM:DATA

## Syntax

```
WFM:DATA <waveform name> <revision> <metadata> <data length>  
    <waveform data>  
DATA? <waveform name> <revision>
```

## Description

Obtain the sample data of a waveform

## Parameters

<waveform name>	Name of the waveform of interest
<revision>	Revision of interest
<data length>	Waveform dimensions (see below)
<waveform data>	Binary encoded waveform data
<metadata>	Waveform metadata

## Notes

- <data length> is the number of dimensions followed by a series of integers, each in square brackets, that define the dimensions of the waveform in samples. For example 3 [65536] [32] [32] specifies that the data is 32x32 waveforms of 65536 points each.
- <waveform data> is binary IEEE floating point (either single or double precision according to the result of WFM:REALSZ) that has been encoded with a binary NOT with (post-inversion) characters 0-32, ';', and '%' replaced with escape sequences. The escape sequence is initiated by the '%' character and consists of '%' followed by the escaped character + 0x80.

## See also

- WFM:REALSZ (pg. 64)
- WFM:METADATA (pg. 63)

# WFM:DATASHM

## Syntax

```
WFM:DATASHM <waveform name> <revision> <metadata> <data
length> <posix shm name>
DATASHM? <waveform name> <revision>
```

## Description

Obtain the sample data of a waveform through POSIX shared memory

## Parameters

<waveform name>	Name of the waveform of interest
<revision>	Revision of interest
<metadata>	Waveform metadata (see WFM:METADATA).
<data length>	Waveform dimensions (see below)
<posix shm name>	Specification of the POSIX shared memory name for access to the binary waveform data.

## Notes

- Query only. Command syntax indicates format of response.
- <data length> is the number of dimensions (an integer) followed by a series of integers, each in square brackets, that define the dimensions of the waveform in samples. For example 3 [65536] [32] [32] specifies that the data is 32x32 waveforms of 65536 points each.
- <posix shm name> is the name of a POSIX shared memory area that can be passed to shm\_open() to obtain access to the waveform data. The data itself is stored as binary IEEE floating point (either single or double precision according to the result of WFM:REALSZ)

## See also

- WFM:REALSZ (pg. 64)
- WFM:DATA (pg. 51)
- WFM:METADATA (pg. 63)

# WFM:DELETEALL

## Syntax

WFM:DELETEALL

## Description

Delete all user-uploaded or user-copied waveforms from the waveform memory.

## Parameters

## Notes

## See also

- [WFM:DELETE](#) (pg. 54)
- [WFM:COPY](#) (pg. 50)
- [WFM:DATA](#) (pg. 51)

# WFM:DELETE

## Syntax

WFM:DELETE <waveform name>

## Description

Delete the specified waveform.

## Parameters

<waveform name>      Name of the waveform

## Notes

- Any revisions of the waveform that are locked will remain in memory, but will not be shown by WFM:LIST or WFM:LISTLOCK.
- Only waveforms created by the user can be deleted with the WFM:DELETE command.

## See also

- WFM:DELETEALL (pg. 53)
- WFM:COPY (pg. 50)
- WFM:DATA (pg. 51)

# WFM:GLOBALREADYREV

## Syntax

```
WFM:GLOBALREADYREV <global revision>  
GLOBALREADYREV?
```

## Description

Wait for a specific global revision to be ready or query the latest ready global revision.

## Parameters

<global revision>      The global waveform revision

## Notes

## See also

- WFM:GLOBALREV (pg. 57)
- WFM:GLOBALREADYREVTIMEOUT (pg. 56)



# WFM:GLOBALREADYREVTIMEOUT

## Syntax

WFM:GLOBALREADYREVTIMEOUT <global revision> <timeout>

## Description

Wait for a specific global revision to become ready with a timeout (optional units, default ms)

## Parameters

<global revision>	The global waveform revision
<timeout>	The timeout, in ms unless otherwise specified

## Notes

- The specified global revision is not locked into memory so it may have been discarded in favor of a more recent (ready) waveform by the time you manage to read it.

## See also

- WFM:GLOBALREVTIMEOUT (pg. 58)
- WFM:GLOBALREADYREV (pg. 55)

# WFM:GLOBALREV

## Syntax

```
WFM:GLOBALREV <global revision>  
GLOBALREV?
```

## Description

Wait for a specific global revision or query the current global revision.

## Parameters

<global revision>      The global waveform revision

## Notes

## See also

- WFM:GLOBALREVTIMEOUT (pg. 58)
- WFM:GLOBALREADYREV (pg. 55)

# WFM:GLOBALREVTIMEOUT

## Syntax

WFM:GLOBALREVTIMEOUT <global revision> <timeout>

## Description

Wait for a specific global revision with a timeout (optional units, default ms)

## Parameters

<global revision>	The global waveform revision
<timeout>	The timeout, in ms unless otherwise specified

## Notes

## See also

- WFM:GLOBALREV (pg. 57)
- WFM:GLOBALREADYREVTIMEOUT (pg. 56)

# WFM:LIST

## Syntax

```
WFM:LIST <waveform count> <global revision> <waveform1 name>  
        <waveform1 revision> <waveform2 name> <waveform2  
        revision> ...  
LIST?
```

## Description

List the current revisions of all waveforms

## Parameters

<waveform count>	Number of waveform descriptions to follow
<global revision>	Global revision count corresponding to this waveform revision set.
<waveform1 name>	Name of the first waveform
<waveform1 revision>	Revision of first waveform

## Notes

- Query only. Command syntax indicates format of response.
- Since this routine does not lock the waveforms into memory there is no guarantee that the specified waveforms or revisions will remain in memory

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:LISTLOCK (pg. 61)
- WFM:LISTREADY (pg. 60)

# WFM:LISTREADY

## Syntax

```
WFM:LISTREADY <waveform count> <global revision> <waveform1  
name> <waveform1 revision> <waveform2 name> <waveform2  
revision> ...  
LISTREADY?
```

## Description

List the current “ready” revision state of all waveforms

## Parameters

<waveform count>	Number of waveform descriptions to follow
<global revision>	Global revision count corresponding to this waveform revision set.
<waveform1 name>	Name of the first waveform
<waveform1 revision>	Revision of first waveform

## Notes

- Query only. Command syntax indicates format of response.
- Since this routine does not lock the waveforms into memory there is no guarantee that the specified waveforms or revisions will remain in memory
- The current “ready” revision state corresponds to a consistent set of waveforms for which all calculations have been completed.

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:LISTREADYLOCK (pg. 62)
- WFM:LISTREADY (pg. 60)

# WFM:LISTLOCK

## Syntax

```
WFM:LISTLOCK <waveform1      name> <waveform1      revision>
              <waveform2 name> <waveform2 revision> ...
LISTLOCK?
```

## Description

List the current revisions of all waveforms and lock those revisions in memory

## Parameters

<waveform1 name>	Name of the first waveform
<waveform1 revision>	Revision of first waveform

## Notes

- Query only. Command syntax indicates format of response.
- You must call WFM:UNLOCK on each of the waveform/revision combinations returned, lest they be stuck in memory. A dropped TCP/IP connection automaticall unlocks all waveforms locked by that connection.

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:UNLOCK (pg. 68)
- WFM:LIST (pg. 59)

# WFM:LISTREADYLOCK

## Syntax

```
WFM:LISTREADYLOCK <waveform1 name> <waveform1 revision>  
  <waveform2 name> <waveform2 revision> ...  
LISTREADYLOCK?
```

## Description

List the current “ready” revision state of all waveforms and lock those revisions in memory

## Parameters

<waveform1 name>	Name of the first waveform
<waveform1 revision>	Revision of first waveform

## Notes

- Query only. Command syntax indicates format of response.
- The current “ready” revision state corresponds to a consistent set of waveforms for which all calculations have been completed.
- You must call WFM:UNLOCK on each of the waveform/revision combinations returned, lest they be stuck in memory. A dropped TCP/IP connection automaticall unlocks all waveforms locked by that connection.

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:UNLOCK (pg. 68)
- WFM:LISTREADY (pg. 60)
- WFM:LISTLOCK (pg. 61)

# WFM:METADATA

## Syntax

WFM:METADATA? <waveform name> <revision>

## Description

Obtain the metadata for a waveform

## Parameters

<waveform name>	Name of the waveform of interest
<revision>	Revision of interest

## Notes

- The response is of the form WFM:METADATA <waveform name> <revision> { <metadatum name>:<type>=<value> <metadatum name>:<type>=<value> ... } <data length>
- <data length> is the number of dimensions (an integer) followed by a series of integers, each in square brackets, that define the dimensions of the waveform in samples. For example 3 [65536] [32] [32] specifies that the data is 32x32 waveforms of 65536 points each.
- Valid types are:
  - integer: <value> is an integer.
  - string: <value> is a quoted string
  - real: <value> is a floating point number.

## See also



# WFM:REALSZ

## Syntax

```
WFM:REALSZ <bytes per floating point number>  
REALSZ?
```

## Description

Returns the number of bytes per floating point number returned by WFM:DATA or WFM:DATASHM.

## Parameters

<bytes per floating point number>      Number of bytes per floating point number.

## Notes

- Query only. Command syntax indicates format of response.
- Should be either 4 (IEEE single precision) or 8 (IEEE double precision).

## See also

- WFM:DATA (pg. 51)
- WFM:DATASHM (pg. 52)

# WFM:REVISION

## Syntax

```
WFM:REVISION <waveform name> <waveform revision>  
REVISION? <waveform name>
```

## Description

List the current revision of the specified waveform

## Parameters

<waveform name>	Name of the waveform
<waveform revision>	Revision of the waveform

## Notes

- Query only. Command syntax indicates format of response.
- Since this routine does not lock the waveforms into memory there is no guarantee that the specified waveforms or revisions will remain in memory

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:LIST (pg. 59)
- WFM:REVISIONLOCK (pg. 66)

# WFM:REVISIONLOCK

## Syntax

```
WFM:REVISIONLOCK <waveform name> <waveform revision>  
REVISIONLOCK? <waveform name>
```

## Description

List the current revision of the specified waveform and lock that revision in memory, or wait for at least a specific revision and lock it in memory

## Parameters

<waveform name>	Name of the waveform
<waveform revision>	Revision of the waveform

## Notes

- Query format finds the latest (not necessarily ready) revision of the specified waveform. The command format waits for the specified revision (or a later version) to become available.
- In both cases the name and revision of the locked waveform are returned.
- You must call WFM:UNLOCK on the waveform/revision combinations returned, lest it be stuck in memory. A dropped TCP/IP connection automaticall unlocks all waveforms locked in memory by that connection.
- This attempts to obtain the specified revision, but may return a more recent version than specified.

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:LIST (pg. 59)
- WFM:REVISION (pg. 65)

# WFM:REVISIONREADYLOCK

## Syntax

WFM:REVISIONREADYLOCK <waveform name> <waveform revision>

## Description

Wait for at least a specific revision of a waveform to become ready, and lock it in memory

## Parameters

<waveform name>	Name of the waveform
<waveform revision>	Revision of the waveform

## Notes

- The command format waits for the specified revision (or a later version) to become available and ready.
- In both cases the name and revision of the locked waveform are returned.
- You must call WFM:UNLOCK on the waveform/revision combinations returned, lest it be stuck in memory. A dropped TCP/IP connection automaticall unlocks all waveforms locked in memory by that connection.
- This attempts to obtain the specified revision, but may return a more recent version than specified.

## See also

- WFM:METADATA (pg. 63)
- WFM:DATA (pg. 51)
- WFM:LIST (pg. 59)
- WFM:REVISION (pg. 65)

# WFM:UNLOCK

## Syntax

WFM:UNLOCK <waveform name> <waveform revision>

## Description

Unlock the specified revision of the specified waveform.

## Parameters

<waveform name>	Name of the waveform
<waveform revision>	Revision of the waveform

## Notes

- The specified waveform and revision must have been previously locked in memory by WFM:REVISIONLOCK or WFM:LISTLOCK.

## See also

- WFM:LISTLOCK (pg. 61)
- WFM:REVISIONLOCK (pg. 66)

# WFM:WFMS

## Syntax

WFM:WFMS?

## Description

Download the full set of user-defined waveforms.

## Parameters

## Notes

- The response begins with WFM:DELETEALL; and continues with semicolon-delimited WFM:DATA commands containing the data and metadata.
- The output is intended to be fed back in to recreate the waveforms

## See also

- WFM:DATA (pg. 51)

## Appendix A

# Dataguzzler native binary file format

The dataguzzler native file format is that written by the dg\_grab download program and that read by the dg\_upload program. The dataguzzler native file format consists of a header followed by a series of chunks. The header is the 8 bytes DATAGUZZ (big endian architectures) or ZZUGATAD (little endian architectures). (recall that all Intel-architecture PC's are little endian). The next 16 bytes is the 8 byte header of the first chunk followed by the 8 byte integer length of the first chunk (not including the length of its header). The first chunk data follows its header. The header for the second chunk follows, etc. Note that all chunks are implicitly padded to be multiples of 64 bits (8 bytes) in size. That is, the written size is the actual size, but zeros are added after the chunk to make an even multiple of 8 bytes before the next chunk is written. Note that on little endian architectures the chunk headers specified below are reversed.

Chunk header	contents
DATAGUZZ GUZZNWFM	wrapper around entire file. Data describing a named waveform. Should contain a WAVENAME chunk followed by a GUZZWFMD chunk
WAVENAME GUZZWFMD	String containing the name of the waveform. Data describing a single waveform. Should contain a METADATA chunk followed by a WFMDIMNS chunk followed by a DATARRAYF or DATARRAYD chunk
METADATA METDATUM	a series of METDATUM subchunks. a METDNAME subchunk containing the metadatum name followed by a METDINTV, METDSTRV, or METDDBLV subchunk
METDINTV METDSTRV	Integer metadatum value, 64 bit signed integer (int64_t) String metadatum value. The size of the chunk is the length of the string
METDDBLV WFMDIMNS	Double precision metadatum value, type double Waveform dimensions. Starts with two 64 bit integers (uint64_t): The product of the dimensions and the number of the dimensions (ndim). This is followed by ndim uint64_t's containing the lengths of the individual dimensions
DATARRAYF	single-precision array data. The first index changes most rapidly
DATARRAYD	double-precision array data. The first index changes most rapidly
SNAPSHOT	Snapshot of an experiment (results for a specific set of parameters). This consists of a METADATA chunk with the parameters of the experiment, followed by a series of GUZZNWFM chunks with the data from the snapshot
SNAPSHTS VIBRDATA	series of SNAPSHOT chunks SNAPSHOT chunk containing configuration parameters and results from vibration measurement: Two SNAPSHOT chunks followed by a VIBFCETS chunk. See boundary_collect_procedure.pdf for more information.
VIBFCETS VIBFACET	One or more VIBFACET chunks Parameters/results of a facet: SNAPSHOT, followed by a series of GUZZNWFM chunks. See boundary_collect_procedure.pdf for more information.



## A.1 Standardized file name extensions

Extension	MIME type	contents
dgz	application/x-dataguzzler-waveform	Single unnamed waveform (GUZZWFMD)
dga	application/x-dataguzzler-array	Array of unnamed waveforms (series of GUZZWFMD chunks)
dgs	application/x-dataguzzler-snapshot	Snapshot of named waveforms (SNAPSHOT chunk, typically with the METADATA section empty) (older version was a series GUZZNWFMD chunks)
dgd set	application/x-dataguzzler-data application/x-dataguzzler-settings	Data from a series of experiments (SNAPSHTS chunk) This is not a chunked format. It is a raw dump of the output of WFM:WFMS? followed by the output of SET?
vibr	application/x-dataguzzler-vibration	Data from a series of vibration measurements (VIBRDATA chunk)

## A.2 File access API

Dataguzzler file I/O has been implemented in libraries written in C, MATLAB/Octave, and Python+Numpy. The bulk of this API documentation describes the C library. The API is similar in the other languages.

### A.2.1 Reading a Dataguzzler file

Open a Dataguzzler file for reading with the `dgf_open()` function:

```
struct dgf_file *infile;  
  
infile=dgf_open(char *filename);
```

`dgf_open()` returns an opaque file handle or NULL if the file could not be opened. After opening the file, the first step is reading the first chunk header. This is done with `dgf_checknextchunk()` or `dgf_nextchunk()` depending on whether or not you know for certain the type of the first chunk.

```
struct dgf_Chunk *Chunk;  
  
Chunk=dgf_nextchunk(struct dgf_file *infile);  
Chunk=dgf_checknextchunk(struct dgf_file *infile, char *chunkname);
```

`dgf_nextchunk()` opens the next available chunk, or returns NULL if there are no chunks left. `dgf_checknextchunk()` opens the next available chunk if `chunkname` matches its eight character name. Otherwise it skips the chunk and returns NULL. If there are no chunks left it also returns NULL.

The `dgf_Chunk` structure contains several useful members: `Name`, `ChunkLen`, and `ChunkPos`

```
struct dgf_Chunk { /* on ChunkStack */
    struct dgl_Node Node;
    char Name[9]; /* 8 characters + 0 terminator so we can use strcmp() et al. */
    int64_t ChunkStart;
    int64_t ChunkLen; /* used only in read routines */
    int64_t ChunkPos; /* relative to ChunkStart */
};
```

- `Name` is a null-terminated copy of the 8-character chunk name string.
- `ChunkLen` is the length of the chunk contents, not including any headers and/or padding.
- `ChunkPos` is the number of bytes that have been read from the chunk.

The chunk can contain either binary data or nested chunks. `dgf_readdata()` can be used to read binary data or `dgf_nextchunk()` can be used to open nested chunks. When done reading the contents of the chunk, you must call `dgf_chunkdone()` to close the chunk. Once the last chunk is closed, call `dgf_close()` to close the file.

```
void dgf_readdata(struct dgf_file *infile, void *Buf, int64_t nbytes);
void dgf_chunkdone(struct dgf_file *infile, struct dgf_Chunk *Chunk); /* 2nd parameter may be NULL */
void dgf_close(struct dgf_file *infile);
```

`dgf_readdata()` reads the specified number of bytes from the current chunk into the specified buffer. `dgf_chunkdone()` closes the most recent open chunk so that following calls to `dgf_<check>nextchunk()` open the following chunk, not nested chunks. `dgf_close()` closes the file after the last chunk is done.

Special-purpose processing routines are included for a few chunk types. These routines should be called immediately after `dgf_<check>nextchunk()` and include calls to `dgf_chunkdone()` to close the chunk.

```
struct dg_wfminfo *dgf_procGUZZWFMD(struct dgf_file *file, char *WfmName);
struct dg_wfminfo *dgf_procGUZZNWFMD(struct dgf_file *file);
void dgf_procMETADATA(struct dgf_file *file, struct dgl_List *MetaData);
```

`dgf_procGUZZWFMD()` processes an unnamed waveform GUZZWFMD chunk. A name (`WfmName`) or NULL may be provided as a parameter. The routine returns a `struct dg_wfminfo *` containing the waveform data. Similarly, `dgf_procGUZZNWFMD()` processes a named waveform GUZZNWFMD chunk, also returning a `struct dg_wfminfo *`. `dgf_procMETADATA()` reads a METADATA chunk, adding the metadata elements to the pre-existing and pre-initialized `MetaData` list.

## A.2.2 Writing a Dataguzzler file

Open a Dataguzzler file for writing with the `dgf_creat()` function:

```
struct dgf_file *outfile;

outfile=dgf_creat(char *Name);
```

`dgf_creat()` returns an opaque file handle or NULL if the file could not be opened. After opening the file, chunks and nested chunks may be written. Create a chunk with `dgf_startchunk()`

```
void dgf_startchunk(struct dgf_file *outfile,char *chunkname);
```

The `chunkname` parameter gives the eight-character null-terminated chunk ID. The chunk can contain nested chunks (create with `dgf_startchunk()`) or binary data (write with `dgf_writedata()`). When done writing the chunk, end it with `dgf_endchunk()`.

```
void dgf_writedata(struct dgf_file *outfile,void *buf,int64_t nbytes);
void dgf_endchunk(struct dgf_file *outfile);
```

When done with the last chunk, close the file with `dgf_close()`.

```
void dgf_close(struct dgf_file *infile);
```

Special-purpose writing routines are included for a few chunk types. These routines start the chunk, write the contents, and end the chunk.

```
void dgf_writenamedwfm(struct dgf_file *file,struct dg_wfminfo *wfm);
void dgf_writewfm(struct dgf_file *file,struct dg_wfminfo *wfm);
void dgf_writemetadata(struct dgf_file *file,struct dgl_List *MetaData);
```

These routines write GUZZNWFM, GUZZWFMD, and METADATA chunks respectively with the provided waveform/metadata.

## A.2.3 MATLAB/Octave file access library

The MATLAB API is essentially similar to the C API. You will need to place the `.m` files somewhere in MATLAB's path. This can be done with the MATLAB `path()` function or with the `MATLABPATH` environment variable. The

primary difference between the MATLAB and C APIs is that since MATLAB parameters are passed exclusively by value rather than by reference, the filehandle structure is processed by each routine and a new filehandle is returned by each routine as an extra result of the function. See `dgf_testread.m` for an example.

#### **A.2.4 Python file access library**

The Python API is essentially similar to the C API. The Numpy (numerical python) library is required.

You will need to `import dg_file`. Then you can access the routines such as `infile=dg_file.dgf_open(filename)`.

#### **A.2.5**