

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note	LIGO-T990013-B - D	6/28/99
<h1>Diagnostics Test Software</h1>		
Daniel Sigg and Peter Fritschel		

Distribution of this draft:

GDS

This is an internal working note
of the LIGO Project.

LIGO Hanford Observatory
P.O. Box 1970 S9-02
Richland, WA 99352
Phone (509) 372-8106
FAX (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (504) 686-3100
FAX (504) 686-7189
E-mail: info@ligo.caltech.edu

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS NW17-161
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: <http://www.ligo.caltech.edu/>

Table of Contents

1 OVERVIEW	5
2 USER INTERFACE	8
2.1 INSTALLATION	8
2.2 GETTING STARTED	8
2.3 COMMAND LINE INTERFACE	9
2.3.1 Diagnostics Tests	9
2.3.2 Arbitrary Waveform Generator	11
2.3.3 Testpoint Control	14
2.4 GRAPHICAL USER INTERFACE	14
2.4.1 Common Properties of Diagnostics Tests	15
2.4.2 Diagnostics Tests	15
2.4.3 Control Screens	15
3 ANALYSIS ALGORITHMS	16
3.1 FFT MEASUREMENTS	16
3.1.1 Sampling rate reduction: multistage decimation	16
3.1.2 Zoom Analysis	19
3.1.3 Data Windowing	21
3.1.4 Power spectrum estimation using the FFT	21
3.1.5 Cross-spectral density	22
3.1.6 Transfer function estimates	23
3.1.7 Coherence estimates	23
3.2 SWEPT SINE MEASUREMENTS	23
3.2.1 Digital demodulation & frequency response calculation	23
3.2.2 Numerical integration algorithm	24
3.2.3 Integration time	26
3.2.4 Settling time	26
3.2.5 Coherence calculation	26
3.3 SINE RESPONSE MEASUREMENTS	26
3.3.1 Harmonic Analysis	26
3.3.2 Two-Tone Intermodulation	27
3.3.3 Transfer Matrix Measurements	27
3.4 POLE-ZERO CURVE FITTING	27
3.5 CORRELATION MEASUREMENTS	27
3.6 TIME MEASUREMENTS	27
3.7 OCTAVE BAND ANALYSIS	27

4 TEST ORGANIZATION	28
4.1 INTERFACES	28
4.1.1 Storage API	28
4.1.2 Real-Time Data Distribution API	29
4.1.3 Excitation API	29
4.1.4 Test Point API	30
4.2 TEST SUPERVISORY	30
4.2.1 Standard Supervisory Task	30
4.3 TEST ITERATORS	30
4.3.1 Repeat	32
4.3.2 Parameter Scan	32
4.3.3 Optimization	32
4.4 TESTS	32
4.4.1 Sine Response	32
4.4.2 Swept Sine	33
4.4.3 FFT Tests	34
4.4.4 Time Series Measurements	36
5 EXCITATION ENGINE	37
5.1 OUTPUT WAVEFORMS	37
5.1.1 Overview	37
5.1.2 Periodic Waveforms	37
5.1.3 Non-Periodic Waveforms	39
APPENDIX A NETWORK INTERFACES	40
A.1 SERVICES INFORMATION	40
A.2 MESSAGE PASSING INTERFACE	41
A.2.1 Using Remote Procedure Calls	41
A.2.2 Using Sockets	42
APPENDIX B DATA REPRESENTATION	44
B.1 COMMON	46
B.1.1 Globals	46
B.1.2 Synchronization Tools	47
B.1.3 Environment	47
B.1.4 Parameter Scan	48
B.1.5 Parameter Optimization	49
B.2 RESULTS	50
B.2.1 Index	50
B.2.2 Plot	50
B.2.3 Time Series	54
B.2.4 FFT and (Cross) Power Spectrum	56
B.2.5 Transfer Function and Coherence	57
B.2.6 List of Coefficients	58
B.2.7 Measurement Values	59

B.3	DIAGNOSTICS TESTS	60
B.3.1	Sine Response, Harmonic Distortion and Two-Tone Intermodulation Tests	60
B.3.2	Swept Sine Tests	61
B.3.3	Fourier Tests	63
B.3.4	Time Series Measurements and Trigger Response Tests	65
B.3.5	Random Stimulus Response Tests	67
B.4	XML CONVENTIONS	67
APPENDIX C	SOFTWARE MODULES	68
C.1	OVERALL STRUCTURE	68
C.2	UTILITIES	70
C.2.1	gdserr	70
C.2.2	gdsheartbeat	70
C.2.3	gdsmutex	70
C.2.4	gdsprm	70
C.2.5	gdsstring	71
C.2.6	gdstask	71
C.2.7	rpcinc	71
C.2.8	tconv	72
C.3	ALGORITHMS	72
C.3.1	decimate	72
C.3.2	gdsrand	72
C.3.3	gdssigproc	72
C.3.4	sineanalyze	72
C.4	ARBITRARY WAVEFORM GENERATOR	73
C.4.1	awg	73
C.4.2	awgapi	73
C.4.3	awgfunc	74
C.4.4	awg_server	74
C.4.5	excitation	74
C.5	COMMAND LINE INTERFACE	74
C.5.1	cmdline	74
C.5.2	gdscmd	75
C.5.3	gdsmsg	75
C.5.4	gdsmsg_server	75
C.5.5	gdsmsg_sockets	75
C.6	DATA ACQUISITION INTERFACE	75
C.6.1	gdschannel	75
C.6.2	gdsrtdd	76
C.7	DIAGNOSTICS TEST UTILITIES	76
C.7.1	diagnames	76
C.7.2	diagorg	76
C.7.3	testenv	76
C.7.4	testsync	76

C.8 DIAGNOSTICS TEST SUPERVISORY	76
C.8.1 supervisory	76
C.8.2 stdsuper	76
C.9 DIAGNOSTICS TEST ITERATORS	76
C.9.1 testiter	76
C.9.2 repeat	76
C.9.3 scaniter	76
C.9.4 finder	76
C.10 DIAGNOSTICS TESTS	76
C.10.1 diagtest	76
C.10.2 sineresponse	78
C.10.3 sweptsine	78
C.10.4 fourier	78
C.10.5 timeseries	78
C.10.6 pseudoran	78
C.11 HARDWARE DRIVERS	78
C.11.1 cobox	78
C.11.2 ds340	78
C.11.3 gsdac	79
C.11.4 gpsclk	79
C.11.5 rmap	79
C.12 REFLECTIVE MEMORY ORGANIZATION	80
C.12.1 testpoint	80
C.12.2 testpoint_server	80
C.13 SCHEDULER	80
C.13.1 gdssched	80
C.13.2 gdssched_client	81
C.13.3 gdssched_server	81
C.14 STORAGE OBJECTS	81
C.14.1 gdsdatum	81
C.14.2 diagdatum	83
C.14.3 rtdinput	83
C.15 PROGRAMS	84
C.15.1 chndump	84
C.15.2 diag	84
C.15.3 gdsd	84
C.15.4 libgds.so	84

1 OVERVIEW

The LIGO global diagnostics system (GDS) provides diagnostics test capability for performing stimulus-response tests. Diagnostics tests are divided into five groups: *(i)* sine response tests which include multiple stimulus and multiple response, harmonic distortion and two-tone intermodulation, *(ii)* swept sine response which determines transfer functions, *(iii)* FFT tools which perform power spectrum estimates and cross-correlation measurements, *(iv)* time series measurements which measure the response to a trigger signal, and *(v)* pseudo-random stimulus response tests which utilize wide-bandwidth excitation signals.

The diagnostics test tools provided by GDS are not meant to replace more traditional means to diagnose the instrument—such as stand-alone network and spectrum analyzers—but rather to complement these tools for cases where test points are not readily available or where measuring points are at far distant locations. In particular, digital servo controllers are implementing digital test points which can be used to inject excitation signals and to extract intermediate feedback signals. The diagnostics test system uses the LIGO data acquisition system to collect signals simultaneously from different subsystems. It implements an excitation engine for generating test signals which are synchronized with GPS time and which are provided to both digital and analog subsystems at all major locations.

Diagnostics tests are run from the control room or from any machine which is connected to the control and monitor network. Data from the instrument are obtained through the network data server which in turn gets the data through a reflective memory network which connects to the data collection units (see Fig. 1 and T980026-00). Data collection units can acquire data from analog signals (ADCU), from digital subsystems (DDCU), or from EPICS channels (NDCU).

Test signals are generated by an excitation engine which is connected to the same reflective memory network which provides both read-back and interface to the digital servo systems. The excitation engine also implements digital-to-analog converters which provide test signals to analog subsystems. Excitation engines are available in every building. Additionally, remote-controlled stand-alone signal generators can be used in temporary setups. Digital servo systems implement a test point interface which allows the user to select a finite set of test inputs and test outputs which are then read or written to and from reflective memory, respectively (see also T980020-A).

Two separate user interfaces to the diagnostics test system are provided: *(i)* a command line interface which allows the user to manually adjust test parameters, start a test and save the results; and *(ii)* a graphical user interface with identical capabilities. The results of a diagnostics test can be stored in the LIGO lightweight data format which follows the extensible markup language (XML) specifications.

A schematic overview of the diagnostics test software is shown in Fig. 2. It consists of the following components:

1. a test organizer which runs tests by dividing them into individual measurement steps and schedules them with the excitation engine, the network data server and the analysis process,
2. a command line interface to the test organizer,

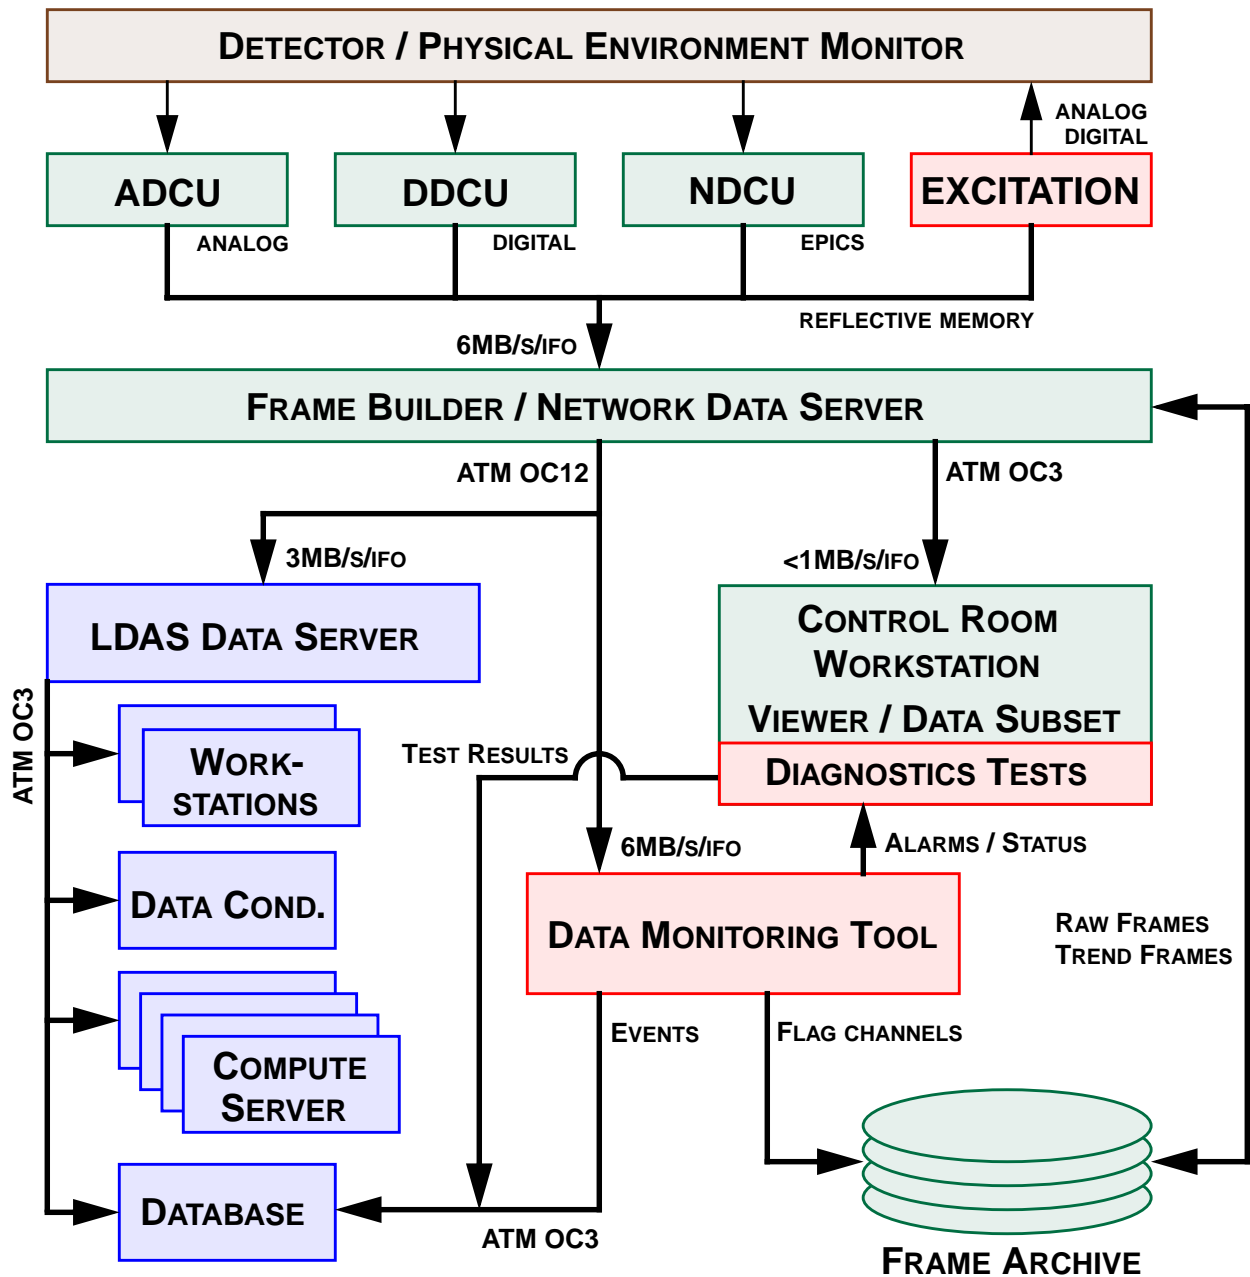


Figure 1: Overview of the data acquisition system (green boxes), the diagnostics system (red) and the data analysis system (blue). The physical environment monitor and the detector are show on top.

3. a graphical user interface for the most commonly used tests,
4. a graphical user interface to the excitation engine for manual operation,
5. a graphical user interface to the testpoint manager,
6. an excitation engine which controls stand-alone function generators (DS340s) and which can generate arbitrary waveforms and write them to a digital-to-analog converter or to digital test points residing in reflective memory,
7. an interface to the network data server to collect data,
8. an algorithm library which is used to analyze the measured data and produce the diagnostics test results,
9. a data wrapper which manages data and parameters associated with a test, and
10. an XML interface which saves and restores diagnostics parameters, data and results to and from disk in the LIGO lightweight data format, respectively.

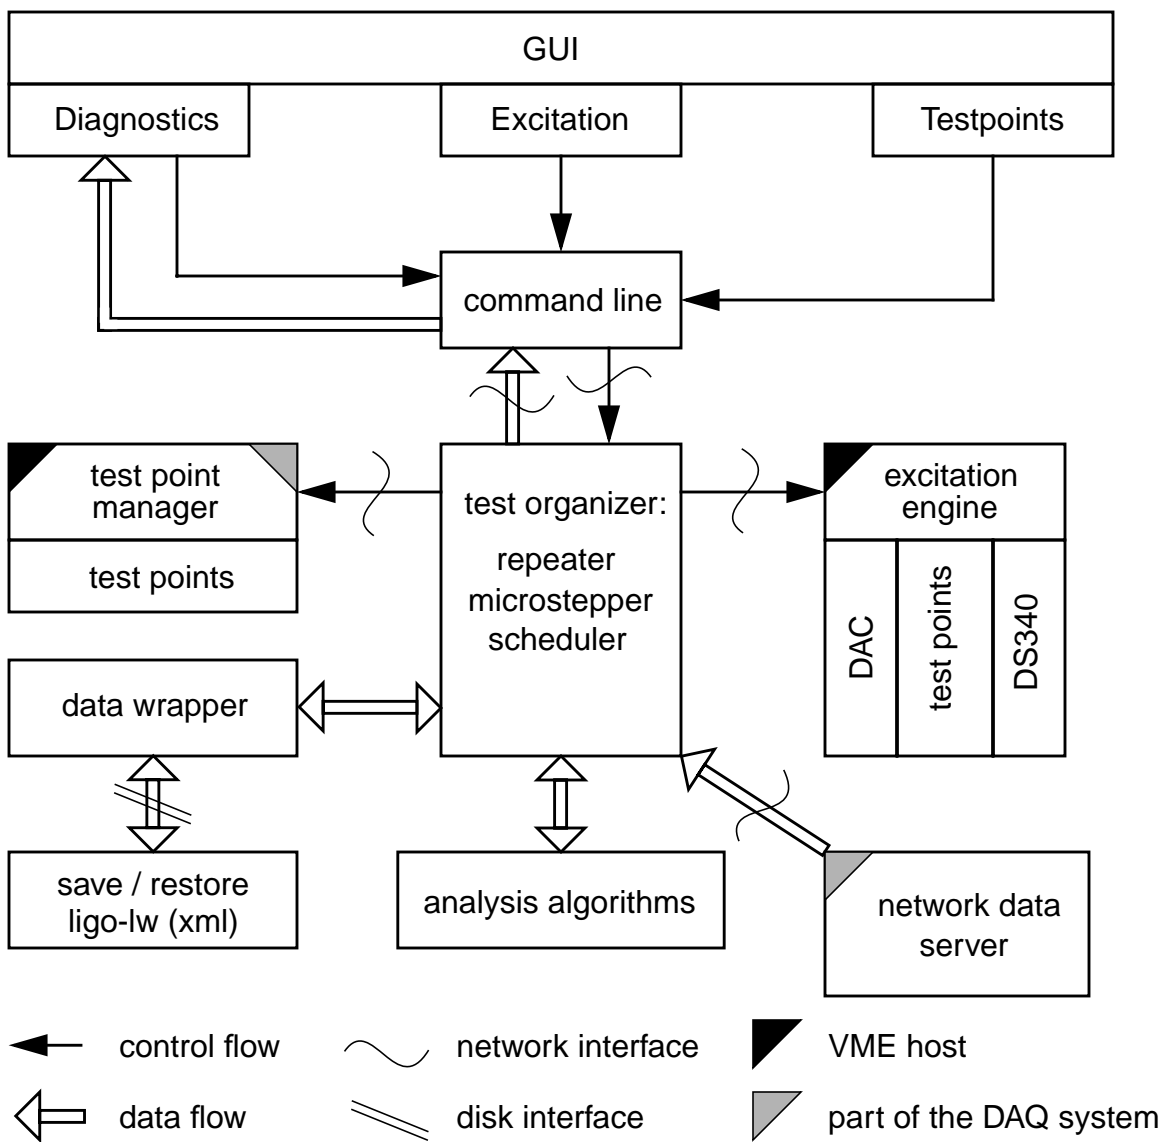


Figure 2: Overview of the diagnostics test software organization.

2 USER INTERFACE

2.1 INSTALLATION

TBD.

2.2 GETTING STARTED

The diagnostics program is started by typing ‘diag –opt’, where ‘–opt’ is a list of options specifying initial start-up parameters. The available options are:

Option	Description
–help	shows a help text.
–g	starts the graphical user interface allowing the user to select a diagnostics test.
–g ‘test’	starts the graphical user interface and preloads the specified test or interface. Possible values for ‘test’ are: awg: arbitrary waveform generator fft: Fourier tests, includes power spectrum and cross-correlation randomresponse: pseudorandom response test, pseudorandom cross spectrum test sineresponse: (multiple) sine response test, harmonic distortion test, two-tone intermodulation test sweptsine: swept sine test timeseries: time series measurement testpoint: testpoint control triggerresponse: single and multiple trigger response test
–c	starts the command line interface of a diagnostics test.
–i	shows information about the configuration and services available.
–l	runs the diagnostics kernel on a the local machine (default).
–s ‘server’	runs the diagnostics kernel on the specified remote machine. Requires that the diagnostics kernel daemon is installed on the remote machine.
–r ‘filename’	reads the specified file as a command line script for initialization. All commands which are available through the command line interface can be used in a script file. Lines starting with a ‘#’ are comments. This script is executed after all other options are parsed, but before the first user input.
–t ‘filename’	uses the specified file as a template for a new diagnostics test.
‘filename’	if exists, opens the specified diagnostics test file; if not, uses it as the default filename for saving the diagnostics test results.

The diagnostics kernel can either run on a local machine or on a remote machine. In the first case it is loaded dynamically and requires the shared object library ‘libgds.so.1’ to be available either in the local directory or the default system library directory.

2.3 COMMAND LINE INTERFACE

The command line interface presents a common way to control diagnostics tests, the arbitrary waveform generator and the test point interface. All these functions are provided by connecting the user interface to a diagnostics kernel, and by passing messages forth and back. The following commands are implemented as part of the user interface:

Command	Description
help	shows help text.
open	loads a diagnostics kernel on the local machine and established a connection. It returns a list of the supported capabilities. Return: supported capabilities: testing testpoints awg Possible error messages: – diagnostics kernel already connected – unable to connect to local diagnostics kernel
open 'server'	opens a new connection to a diagnostics kernel on the specified remote machine. Return: Possible error messages: – diagnostics kernel already connected – unable to connect to remote diagnostics kernel
close	closes the connection to the diagnostics kernel Return: none Possible error messages: – not connected to a diagnostics kernel
read 'filename'	opens the specified file and interprets each line of the file a command. Lines which start with '#' are ignored. Return: echo of commands found in file Possible error messages: – file not found
exit/quit	quits the command line interface after closing any open connections.

Any command which is not recognized by the command line interface will be send to the diagnostics kernel. In case of an error the message “unable to send command to diagnostics kernel” is returned, otherwise the reply string of the command is displayed.

2.3.1 Diagnostics Tests

Diagnostics tests are controlled through named variables. A named variable can contain almost anything, for instance, it can contain a test parameter, a time series of the raw data, or a result array. Thus, only a small set of commands is necessary to setup and run a diagnostics test; they can be separated into the following categories:

- i)* commands to set and read parameters,
- ii)* commands to start and stop a diagnostics test,
- iii)* commands to save or restore settings, data and results of a diagnostics test, and
- iv)* notification messages which are passed back from the diagnostics kernel to inform the caller that a test has finished, or that a new result is available.

When running the diagnostics kernel on a remote machine the save and restore functions could in principle access either the local or the remote file system. The command line interface supports both possibilities, but defaults to the local machine.

The following list of commands is used to control diagnostics tests:

Command	Description
help 'test'	shows help text for specified test. Use "help *" for a list of available tests.
run	starts a diagnostics test. Return: running Possible error messages: – –
abort	aborts a diagnostics test. Return: aborted Possible error messages: – test is not running
pause	Pauses a diagnostics test. Return: test paused Possible error messages: – test is not running
resume	Resumes a paused diagnostics test. Return: test resumed – test is not running
set 'variable name' = 'value'	Sets the specified variable to the given value(s). Return: none Possible error messages: – illegal argument
get 'variable name'	Gets the value(s) of the specified variable(s). The variable name can contain the wildcard character "*". Return: variable(s) and value(s) Possible error messages: – illegal argument
del 'object name'	Erases a data objects belonging to scans, environment, plot, channel data and result. Return: 'object' erased Possible error messages: – illegal argument
defined 'variable name'	Tests if a variable of the given name is defined. Use "defined *" to print a list of all defined variables. Return: yes/no, or list of names Possible error messages: none
brief	if on (default), the value returned by get is truncated if it is very long. Return: brief is on/off. Possible error messages: none

Command	Description
save –'flags' 'filename'	saves the data of a diagnostics test. The flag specifies the data to be saved: all (all data associated with a diagnostics test), ext (extended data, all but images), std (standard data: parameters, settings and result), or prm (parameter data only). If the flag is omitted the default behavior is standard. If the filename starts with a ':' and if the command line interface is connected to a remote machine, the test is saved on the remote machine. Return: 'filename' saved Possible error messages: – illegal filename
restore –'flags' 'filename'	restores the data of a diagnostics test. The flag specifies the data to be restored: all (all data associated with a diagnostics test), ext (extended data, all but images), std (standard data: parameters, settings and result), or prm (parameter data only). If the flag is omitted the default behavior is standard. If the filename starts with a ':' and if the command line interface is connected to a remote machine, the test is restored from the remote machine. Return: 'filename' restored Possible error messages: – file not found

Each diagnostics test defines a set of valid named variables—including their data type—which are necessary to run the test, or which are returned as a result. Variable names should not contain spaces, tabs, nor any other special character. Variable names are case-insensitive.

2.3.2 Arbitrary Waveform Generator

Arbitrary waveform generators are implemented separately for every interferometer node. An interferometer node can have multiple excitation engines (CPUs or stand-alone signal generators), each of them consisting of multiple slots which can be used independently to send waveforms to different channels. Commands to the arbitrary waveform generator are of all the form 'awg command arguments'. The following commands are supported:

Command	Description
help	shows help text.
channels	displays all excitation channel names
show 'node':awg'	shows the configuration of an arbitrary waveform generator. Return: AWG information Possible error messages: – arguments for show are 'node'.awg' – node%i/awg %i not available

Command	Description
new 'channel'	reserves a slot in an arbitrary waveform generator for the specified channel. Requires the full channel name and returns a non-negative slot number if successful. (The returned slot number also encodes 'node' and 'awg' number and is unique for each system.) Return: slot %i Possible error messages: – no slot available or invalid channel name
free 'slot'	frees a previously reserved slot of an arbitrary waveform generator Return: slot %i freed Possible error messages: – invalid slot number – slot not available or invalid
set 'slot' 'waveform' {, 'waveform'}	selects a new waveform in the given slot of an arbitrary waveform generator. If the waveform argument is omitted the slot is cleared and the output reset to zero. Return: slot %i enabled Possible error messages: – invalid slot number – invalid arguments – not enough memory – unable to download waveform – unrecognized waveform – slot not available or invalid More than one waveform can be specified by using a comma separated list of waveforms.
clear 'node'. 'awg'	clears all waveforms from a given arbitrary waveform generator and frees all slots. Return: reset succeeded Possible error messages: – reset failed – node %i/awg %i not available
stat 'node'. 'awg'	shows statistical performance data of a given arbitrary waveform generator. Return: statistics information Possible error messages: – arguments for stat are 'node'. 'awg' – node %i/awg %i not available – no statistics available The statistical record can be reset by specifying an additional 'r'.

The '%' arguments in the return message are replaced by the corresponding value and follow the C printf convention. Error message are preceded by the word "error:". If the command is none of the above, the message "unrecognized command" will be returned.

The command arguments are explained below:

Argument	Description
'node'	interferometer node, e.g. 0 for H1 and L1, and 1 for H2.
'awg'	identification number of excitation engine, or stand-alone signal generator: starts with 0, and the stand-alone units generally use 5.
'channel'	full channel name of an excitation or read-back channel, e.g. "H1:LSC-TEST_1".
'slot'	slot number as returned by new.
'waveform1'	specifies a periodic waveform of format: 'func' 'frequency' 'amplitude' 'offset' 'phase' 'ratio' where the frequency is given in Hz, the amplitude and phase in V, and the phase in rad. The ratio parameter can be used to specify the ratio between high and the full period of a square wave (default is 0.5).
'waveform2'	specifies an impulse: impulse 'frequency' 'amplitude' 'duration', 'delay' A frequency value of 0 indicates a single pulse.
'waveform3'	specifies a constant offset: const 'amplitude' where the amplitude is specified in V.
'waveform4'	specifies a band-limited noise source of format: 'noise' 'start frequency' 'stop frequency' 'amplitude' 'offset'
'waveform5'	specifies a sweeping sine wave of format: sweep 'start frequency' 'stop frequency' 'start amplitude' 'stop amplitude' 'sweep time' 'sweeptype' 'updn' where the sweep time is given in sec.
'waveform6'	specifies an arbitrary waveform of format: arb 'sample frequency' 'scaling' 'trigger type' 'rate' 'value 1' 'value2'... 'value n' where the values are given in V. The output voltage is calculated by multiplying the specified values with the scaling factor. The trigger can either be continuous which repeats the waveform indefinitely, wait which waits for a trigger, or trigger which outputs the (previously stored) waveform exactly once. If a continuous trigger is specified, the rate indicates trigger rate. If a random trigger is specified, the arbitrary waveform will be applied at random time times with an average rate compatible with the specified one.
'func'	one of the following: sine, square, ramp, or triangle.
'noise'	one of the following: normal (normal distributed noise), or uniform (uniformly distributed noise).
'sweeptype'	one of the following: linear, or log.
'updn'	one of the following: + (up), - (down), or 'blank' (bidirectional)
'trigger type'	one of the following: c (continuous), r (random), w (wait), or t (trigger)

2.3.3 Testpoint Control

Each interferometer node has its own testpoint control interface. Commands to the testpoint interface are of all the form ‘tp command arguments’. The following commands are supported:

Command	Description
help	shows help text.
show ‘node’	shows the configuration of a test point interface. Return: testpoint information Possible error messages: – invalid node number supported wildcard: “show *” to show test point information of all nodes.
set ‘node’ ‘number1’... ‘number n’	sets testpoints of the specified node. Return: test point set Possible error messages: – invalid node number – unable to set test point
clear ‘node’ ‘number1’... ‘number n’	Clears testpoints from the specified node. Return: test point cleared Possible error messages: – invalid node number – unable to clear test point supported wildcards: “clear *” to clear all testpoints from all nodes, and “clear ‘node’ *” to clear all testpoints from the specified node.

Error message are preceded by the word “error:”. If the command is not recognized, the error message “unrecognized command” will be returned.

2.4 GRAPHICAL USER INTERFACE

The graphical user interface is using the same message passing interface to communicate with the diagnostics kernel. Thus, all functions of the diagnostics kernel which are accessible through the command line interface can also be used by the graphical user interface. A description of the message passing mechanism can be found in Appendix A.2.

2.4.1 Common Properties of Diagnostics Tests

- 2.4.1.1 File Handling**
- 2.4.1.2 Synchronization Tools**
- 2.4.1.3 Channel Selection**
- 2.4.1.4 Exporting Data**
- 2.4.1.5 Parameter Sweep**
- 2.4.1.6 Parameter Optimization**

2.4.2 Diagnostics Tests

- 2.4.2.1 Sine Response, Harmonic Distortion and Two-Tone Intermodulation Tests**
- 2.4.2.2 Swept Sine Tests**
- 2.4.2.3 Fourier Tests**
- 2.4.2.4 Time Series Measurements**
- 2.4.2.5 Trigger Response Tests**
- 2.4.2.6 Random Stimulus Response Tests**

2.4.3 Control Screens

- 2.4.3.1 Arbitrary Waveform Generator**
- 2.4.3.2 Testpoint Control**

3 ANALYSIS ALGORITHMS

3.1 FFT MEASUREMENTS

An FFT program (such as the FFTW, Fastest Fourier Transform in the West) will be used to compute the N -point DFT (discrete-time fourier transform) of a data stream x :

$$X(k) = \sum_{n=0}^{N-1} x(n) \cdot e^{-2\pi i k n / N} \quad (1)$$

3.1.1 Sampling rate reduction: multistage decimation

Performing an N -point FFT analysis of a data channel at the sampled rate gives a full span analysis, with spectral information in the band $0-f_s/2$. For baseband measurements (frequency spans which start at DC) with increased resolution, the duration of the time record must be increased. This is done by decimating the data appropriately, while keeping the FFT length fixed at N . To avoid aliasing in the decimation process, the data must be low-pass filtered with a digital filter. For large decimation factors, it is more efficient to implement multiple stages of decimation, as shown in Fig. 3.

To simplify (though not necessarily optimize) the computation we fix the single stage decimation at a factor of $D_i = 2$. Thus decimation factors of 2^n (n an integer) are possible by computing n

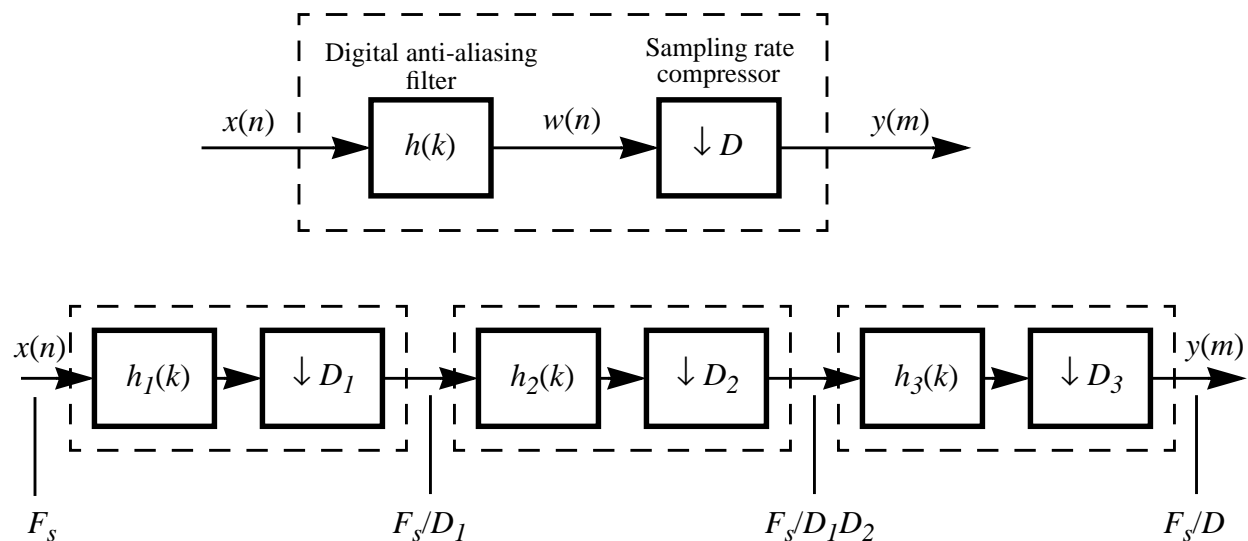


Figure 3: Top: Block diagram of a decimator. The digital low-pass filter is an FIR filter:

$$w(n) = \sum_{k=0}^{N-1} h(k)x(n-k),$$

where N is the number of filter coefficients. The sampling rate compressor simply selects every D th output sample. Bottom: Multistage decimation process.

stages of decimation, leading to FFT spans of $DC - F_s/2^{n+1}$ (where F_s is the channel sampling rate). This simplification means that the coefficients of the digital low-pass filters in each stage are identical: $h_1 = h_2 = \dots = h_n$.

Sampling rate converters generally implement FIR digital filters because of their linear phase response, a necessary feature in our application as well. The diagrams in Fig. 3 imply a direct-form realization of the FIR filter. This is an inefficient method of calculation, since the filter is operating at the full input sample rate, but only every D th output sample is kept. It is more efficient to embed the downsampling operation within the filter, as shown by the realization of Fig. 5. This structure also takes advantage of the symmetry of the filter coefficients for a linear-phase FIR filter – i.e., that $h(n) = h(N - 1 - n)$; this allows halving the number of multiplications to compute the filter.

Another technique to reduce the computation time in a decimate-by-2 stage is to use a *half-band* filter. A half-band filter is one which satisfies two constraints: the ripple in the pass- and stop-bands are equal; the pass- and stop-band cutoff frequencies are related to the Nyquist frequency as: $f_p + f_s = f_N$. These symmetries result in about 50% of the filter coefficients being zero, cutting down the number of multiplies by roughly a factor of 2. Care must be taken in the design of the digital filter so that when multiple stages are used (up to ~ 10 stages may be used in practice), the passband ripple from each stage does not accumulate into a much larger ripple at the output. Two types of filter design methods have been used to generate specific filters:

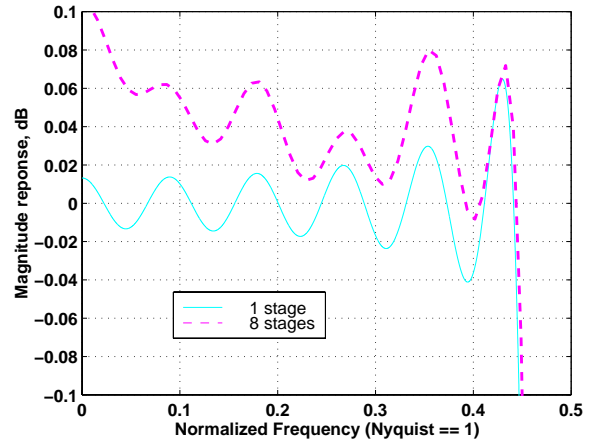
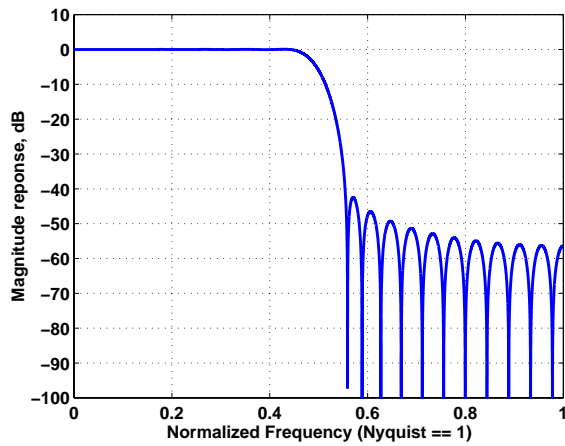
1. An FIR filter designed using a least squares technique trades off increased error at the band edge for better response over most of the passband; thus a cascade of these filters retains small passband error, since the band edge error of one filter is cut off by the next filter.
2. The McClellan-Parks filter design algorithm gives a constant error over the pass- and stop-bands. Compared to the least squares design, this filter has better performance (lower ripple) near the band edge; however, multiple stages build up more error in the lower part of the passband.

Frequency responses for both types of filter, each of order 42, are shown in Fig. 4; each is computed with 12 multiplications per output point. Several FIR filters are available to choose from (more can always be added) to trade-off between execution speed and error. Current selections are:

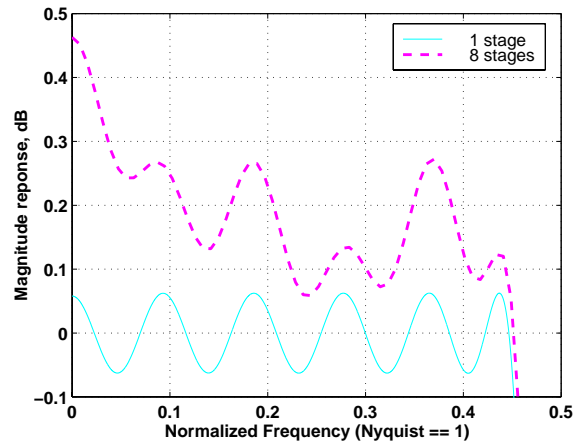
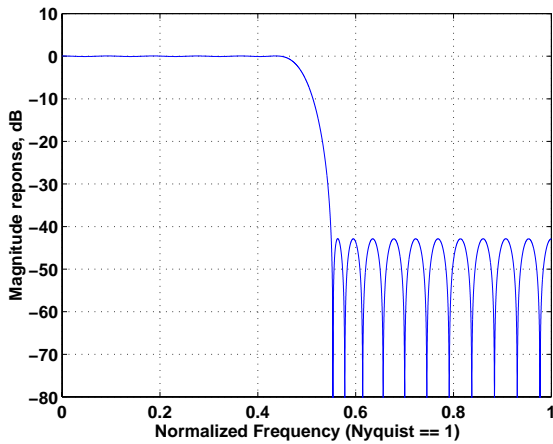
Id	Type	Order	Passband cutoff frequency ¹	Passband Ripple (dB)		Stopband attenuation (dB)
				Midband	Edge	
2	Least-squares	22	0.45	0.1	0.8	30-40
1	Least-squares	42	0.45	0.02	0.1	40 - 56
3	Least-squares	82	0.45	0.0006	0.01	60 - 90
4	McClellan-Parks (equiripple)	42	0.45	0.05		43

Table 1: Decimation filters.

1. as a fraction of the Nyquist frequency



(a)



(b)

Figure 4: Magnitude response of two half-band, order 42 FIR filters. On the left is the response of a single filter stage, and the right shows the passband response for a single stage (solid line) and for 7 cascaded decimation stages (dashed line). (a): filter designed using a least-squares method; (b) filter designed using the McClellan-Parks method. The normalized pass- and stop-band cut-off frequencies are 0.45 and 0.55, respectively.

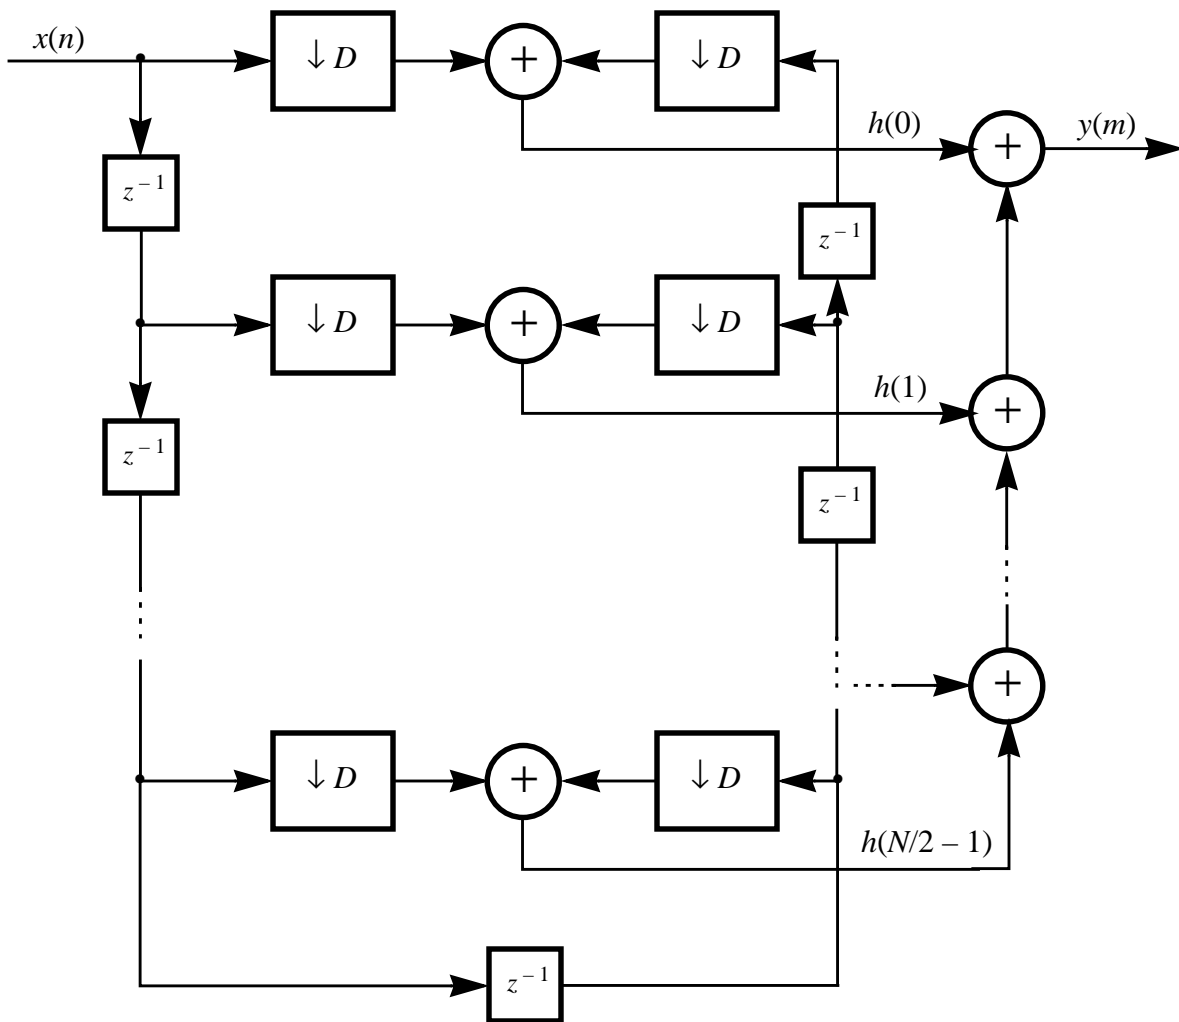


Figure 5: Efficient structure for computation of a decimator that exploits the symmetry of the FIR filter coefficients. For a half-band filter, half of the coefficients h would be zero, and thus not computed.

3.1.2 Zoom Analysis

In order to start an FFT span above DC, the data must first be down-converted so that the center frequency of the span is shifted to DC. This is accomplished with the heterodyne procedure shown in Fig. 6. The heterodyne multiplication produces a complex sequence y from the real data sequence x :

$$y(n) = e^{-j2\pi f_c n} \cdot x(n) \quad (2)$$

The DFT of this sequence,

$$Y(k) = \sum_{n=0}^{N-1} x(n) \cdot e^{-2\pi i \frac{n}{N}(k+f_c)} \quad (3)$$

shows that the frequencies k of the original sequence x have been shifted to the frequencies $k - f_c$ of the sequence y . After heterodyning, we need to low-pass filter and decimate the data to select the frequency span of interest. The data are now complex, so a complex FFT is required; the negative frequencies of Y in this case are independent of the positive frequencies, and in fact they contain the lower half of the span, as shown in Fig. 6.

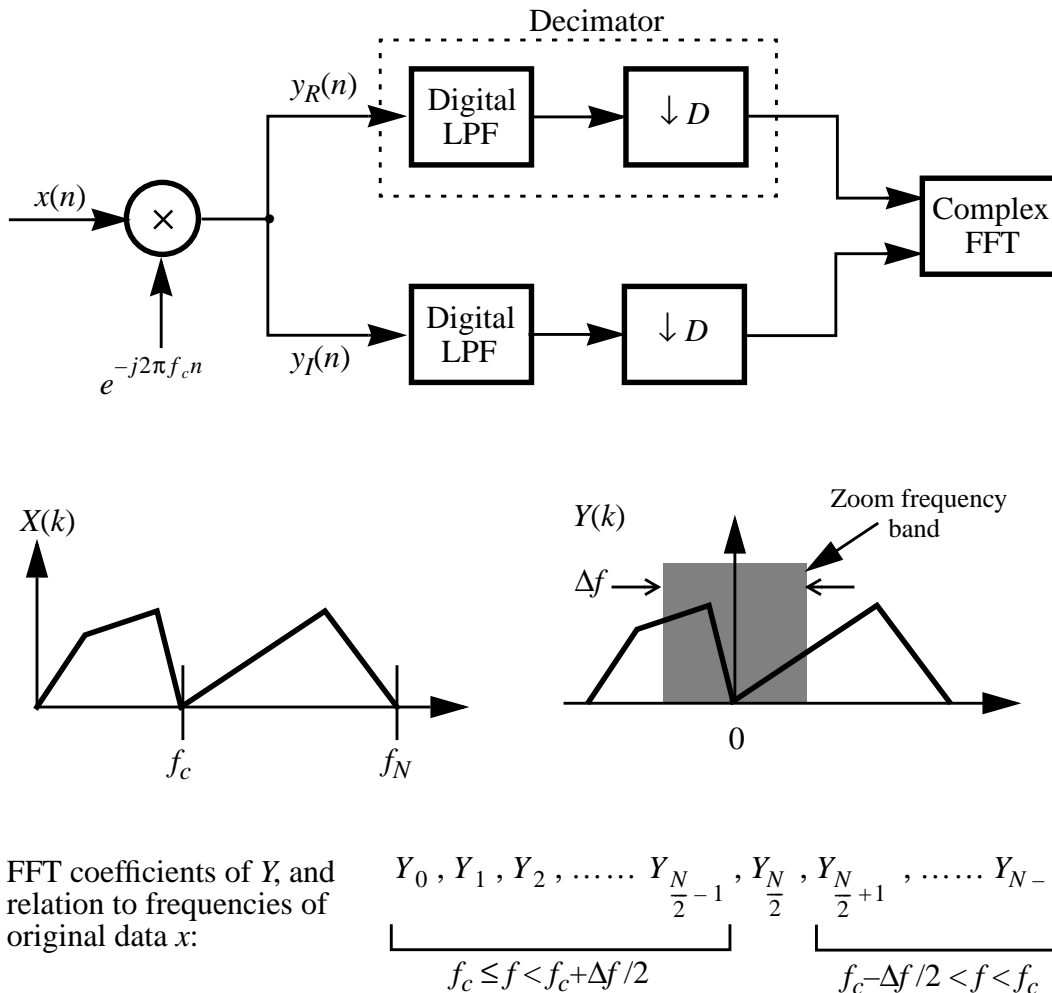


Figure 6: Block diagram of heterodyning for zoom analysis. The center frequency f_c is downconverted to DC, and the lower half of the frequency span is mapped to negative frequencies. The full span of the zoom, Δf , is again restricted to $(1/2^p)$ multiples (p an integer) of the original channel sampling frequency. The decimator consists of $(F_s/2\Delta f)$ stages of the single

3.1.3 Data Windowing

All FFTs are performed on windowed data. The following windowing functions are offered:

Window Type	Time-domain sequence $0 \leq i \leq N-1$	Max. Amplitude error	Comments
Uniform	no window	4 dB	useful for looking at transients
Hanning	$w_i = \frac{1}{2} \left[1 - \cos\left(\frac{2\pi i}{N-1}\right) \right]$	1.5 dB	general purpose; lowest noise floor
Flat-top	$w_i = \frac{1}{2} \left(1 - 1.93 \cos\left(\frac{2\pi i}{N-1}\right) \right)$ $+ 1.29 \cos\left(\frac{4\pi i}{N-1}\right) - 0.388 \cos\left(\frac{6\pi i}{N-1}\right)$ $+ 0.028 \cos\left(\frac{8\pi i}{N-1}\right)$	0.02 dB	very wide pass band, but most accurate amplitude meas.
BMH	$w_i = \frac{1}{2} \left(1 - 1.36109 \cos\left(\frac{2\pi i}{N-1}\right) \right)$ $+ 0.39381 \cos\left(\frac{4\pi i}{N-1}\right)$ $- 0.032557 \cos\left(\frac{6\pi i}{N-1}\right)$	0.8 dB	large dynamic range (good for separating two close frequencies, with widely different amplitudes)

Table 2: Windowing functions that can be applied to the data prior to Fourier transforming. Amplitude error in the transformed data occurs for frequencies not exactly at a bin frequency.

3.1.4 Power spectrum estimation using the FFT

Welch's method of power spectrum estimation is implemented according to the following algorithm:

1. the data is broken up into D , possibly overlapping segments (50% overlap typical), of N points each; each segment may then optionally be 'de-trended' (i.e., the mean, or best linear fit may be removed) before being windowed
2. each (windowed) segment is transformed with an N -point FFT, giving $X_i(k)$ ($i = 1 - D$).

- the magnitude squared of each segment is properly scaled, forming the periodogram $P_i(k)$ of each segment:

$$P_i(0) = \frac{1}{W} |X_i(0)|^2$$

$$P_i(k) = \frac{1}{W} [|X_i(k)|^2 + |X_i(N-k)|^2] \quad k = 1, 2, \dots, \left(\frac{N}{2} - 1\right) \quad (4)$$

$$P_i(N/2) = \frac{1}{W} |X_i(N/2)|^2$$

where W accounts for the power in the window:

$$W \equiv N \sum_{n=0}^{N-1} w^2(n) \quad (5)$$

This normalization is chosen so that the value at each point in the periodogram gives the energy density in units of mean-squared-amplitude.

- the average of the D periodograms is computed to form the power spectrum of x :

$$\langle P_{xx}(k) \rangle = \frac{1}{D} \sum_{i=1}^D P_i(k) \quad (6)$$

- scaling factors may be applied to convert from mean-squared-amplitude (msa) to msa/Hz, (root-mean-square)/Hz^{-1/2}, peak, engineering units, etc.

If the data has been mixed down for zoom analysis, the periodogram is formed somewhat differently, as indicated in Fig. 6.

3.1.5 Cross-spectral density

The cross-spectral density between two data streams $x(n)$ and $y(n)$ (the sampling rates of the two channels, if not already equal, must be equalized by the sample rate conversion described above) is computed as follows:

- each data stream is segmented and de-trended & windowed as above
- each segment of each channel is transformed with an N -point FFT
- the FFT segment pairs are multiplied as:

$$P_{xy}(k) = \frac{1}{W} X(k) \cdot Y(k)^* \quad , \quad k = 0, 1, \dots, N-1 \quad (7)$$

where $*$ denotes complex conjugate.

- the average over the segments is taken, as above, to form $\langle P_{xy}(k) \rangle$

3.1.6 Transfer function estimates

The transfer function estimate given an input channel x and an output channel y is the quotient of the cross-spectrum of x and y and the power spectrum of x :

$$\langle T_{xy}(k) \rangle = \frac{\langle P_{xy}(k) \rangle}{\langle P_{xx}(k) \rangle}, \quad k = 0, 1, \dots, \frac{N}{2} \quad (8)$$

3.1.7 Coherence estimates

The coherence between two signal vectors x and y is computed as:

$$\begin{aligned} \langle C_{xy}(0) \rangle &= \frac{|\langle P_{xy}(0) \rangle|^2}{\langle P_{xx}(0) \rangle \cdot \langle P_{yy}(0) \rangle} \\ \langle C_{xy}(k) \rangle &= \frac{[|\langle P_{xy}(k) \rangle|^2 + |\langle P_{xy}(N-k) \rangle|^2]}{\langle P_{xx}(k) \rangle \cdot \langle P_{yy}(k) \rangle} \quad k = 1, 2, \dots, \left(\frac{N}{2} - 1\right) \\ \langle C_{xy}(N/2) \rangle &= |\langle P_{xy}(N/2) \rangle|^2 \end{aligned} \quad (9)$$

3.2 SWEPT SINE MEASUREMENTS

3.2.1 Digital demodulation & frequency response calculation

Consider a signal $s(t)$ which contains a sine wave of frequency ω , of which we wish to know the amplitude; expressed as a Fourier series:

$$s(t) = \sum_{n=-\infty}^{\infty} c_n \cdot e^{i\omega n t} \quad (10)$$

The desired coefficient is calculated using the integral

$$c_1 = \frac{1}{T} \int_0^T s(t) e^{-i\omega t} dt \quad (11)$$

where $T = 2n\pi/\omega$ (i.e., an integral number of cycles). We only have a sampled version of $s(t)$, of course, which complicates the integration over an integral number of cycles; this is addressed in the next section. The steps in making a sine detection and frequency response calculation are:

1. For the desired channel, obtain the number of data points required by the integration routine (see below).
2. Multiply the data set by $e^{-i\omega t}$.
3. Integrate the real and imaginary parts of (2) over an integral number of sine wave cycles, using the numerical integration algorithm (see below), to compute the complex amplitude c_1 .
4. Repeat steps 1-3 for the desired number of averages (user-specified).

5. The values of c_1 are averaged to form $\langle c_1 \rangle$; this can be used to compute the (complex) frequency response with other channels. The individual coefficients c_1 from each of the average measurements are also stored; these are needed to compute the coherence with other channels (see below).

3.2.2 Numerical integration algorithm

The numerical integration of equation (11) is done with a modified Newton-Cotes method, as follows. We find the fifth-order polynomial that passes through six adjacent points of the demodulated data series (real and imaginary treated separately), and integrate this polynomial over the middle two points. We then move forward one sample, and find and integrate the new polynomial; this process is repeated until the end of the data set, the sum of all the steps giving us the desired integral.

The last integration step must be modified to end the integration at an exact integral number of sine wave cycles. This is done by again finding the fifth-order polynomial that passes through the last six points, but only integrating over a portion of the span covered by the middle two points, stopping where a total span of an integral number of cycles has been reached.

Such an integration step can be reduced to 6 coefficients which multiply the 6 adjacent data points to give the integral over the middle two points; the sum of these 6 coefficients is unity (try integrating a constant function $f(x) = 1$ using this method). Thus, when the whole data set is convolved with these 6 coefficients, nearly all data points are simply multiplied by unity; only 5 points at the beginning and 6 points at the end are multiplied by non-unity coefficients. Note that for each different detection frequency or integration time, the trailing coefficients must be recomputed (the lead coefficients are always the same).

This integration algorithm is similar to Simpson's rule, but uses a higher order approximation for increased accuracy. The accuracy is tested by integrating

$$\frac{1}{T} \int_0^T e^{-in\omega t} dt \quad (12)$$

which should integrate to zero. This is shown in Fig. 7, as a function of the number of data points in a cycle. As this shows, the method loses accuracy as the frequency increases (fewer points per cycle). To further improve the accuracy, we also pass the demodulated data through an FIR low-pass filter, prior to integration. We use a 20-tap FIR filter, with a cut-off frequency of $0.1 \times (\text{Nyquist frequency})$. In fact, the convolution of this filter with the data can be combined with the integration operation, giving a long coefficient vector that is simply dotted (dot-product) with the data. Most of the coefficients in this vector are unity, with 24 non-unity lead coefficients and 25 non-unity trailing coefficients (the latter must be computed for each detection frequency).

The number of data points used by the integration algorithm, per average, is given by

$$n_{\text{pt}} = N_{\text{lead}} + 1 + \text{floor}(t_{\text{int}} \cdot f_s) \quad (13)$$

where $N_{\text{lead}} = 24$ is the number of leading coefficients, t_{int} is the integration time (corresponding to an integral number of cycles of the sine wave), f_s is the sampling frequency, and 'floor(x)' gives the largest integer not greater than x .

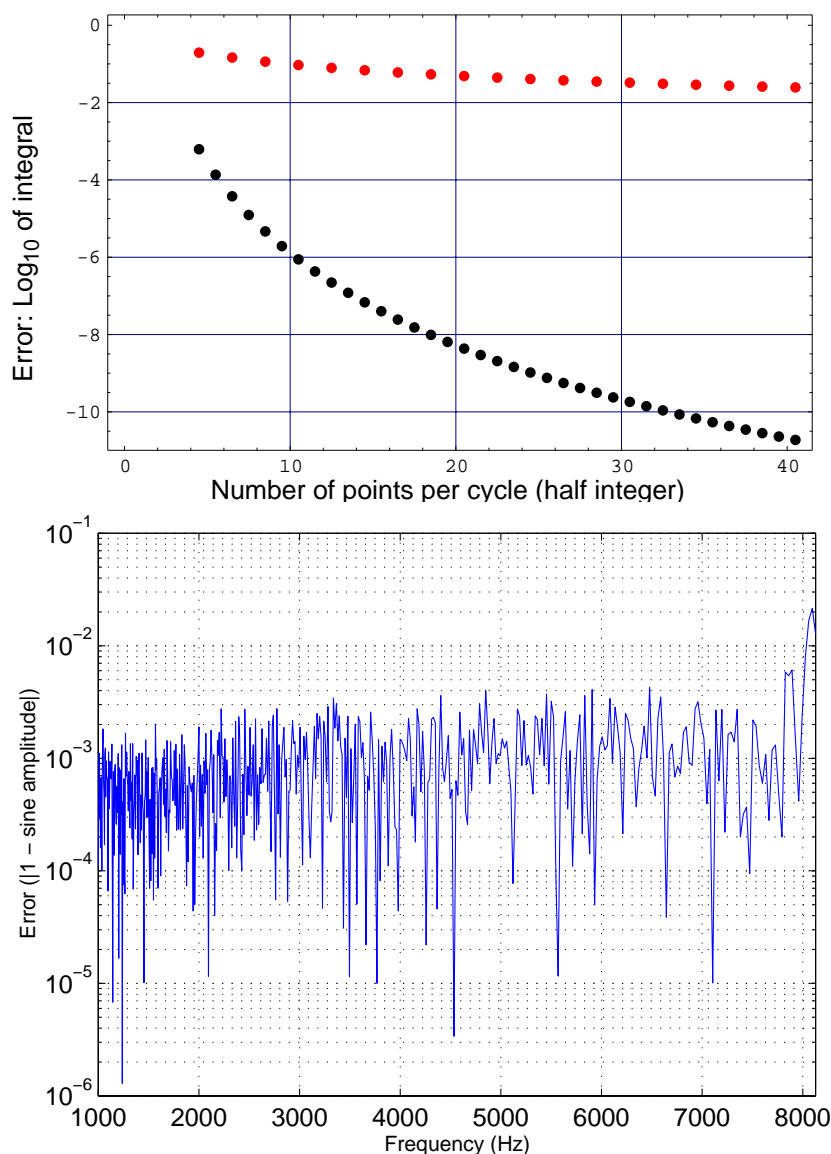


Figure 7: Accuracy of numerical integration algorithm. **Top:** plotted is the integral of equation (12) as a function of the number of points in a cycle, using Simpson's rule (modified to work over a fraction of the last span) (top/red), and the modified 5th-order polynomial method (bottom/black). The points are all computed using a half-integer number of points in a cycle, since this is where the error is the largest. **Bottom:** Error in the integration of a unit amplitude sine wave, plus: a DC offset of 100; uniformly distributed random values in the range 0-0.1. The sine wave is sampled at 16384 Hz, and each frequency is detected with 2 averages of 100 cycles. The error increase above ~ 7.7 kHz because there are very few data points per cycle for the integration algorithm; below 7.7 kHz the error is due to the random values (noise) added to the data.

3.2.3 Integration time

The integration time may be specified either as a time duration, or in the number of sine wave cycles. If specified as a time interval, it is checked that it equals an integral number of cycles of the sine detection frequency; and if not, it is set equal to the nearest cycle number. For detection frequencies near a channel's Nyquist frequency, we want to ensure that a sufficient number of cycles are integrated over, and so we set a minimum integration time that is equal to 10 cycles at the Nyquist frequency (1.2 msec for a 16384/sec sample rate).

3.2.4 Settling time

At each new frequency in a swept sine test (or a sine response test), the test channels are allowed to 'settle' before a sine detection is performed on the data. The settling time is nominally 10% of the integration time. See also section 4.4.

3.2.5 Coherence calculation

Given a pair of channels on which multiple-average sine detection has been performed, the coherence between these two channels can be calculated as follows:

$$\frac{|\langle c_1 \cdot d_1^* \rangle|^2}{\langle |c_1|^2 \rangle \langle |d_1|^2 \rangle} \quad (14)$$

where $\{c_1\}$ and $\{d_1\}$ are the sets of complex amplitudes resulting from n sine detections on each of the two channels, and the averaging is performed over these n values.

3.3 SINE RESPONSE MEASUREMENTS

A sine response test is using the same integration algorithm as a swept sine test to determine transfer coefficients, but can use multiple sine wave excitations simultaneously. This allows to determine non-linearities and transfer matrices of multi-dimensional input-output problems. Since response tests with a single excitation signal will automatically compute a harmonic analysis, tests with exactly two excitations will perform a two-tone intermodulation analysis and tests with two or more excitations and with the same number of excitation readback channels will automatically compute the transfer matrix.

3.3.1 Harmonic Analysis

Assuming an excitation signal is applied at frequency f , and the complex amplitude/phase coefficient, a_f , is measured on a channel b , then the complex amplitude/phase coefficients at integer multiples of the fundamental frequencies represent the harmonic coefficients, $a_{2f}, a_{3f}, \dots, a_{Nf}$. The total harmonic distortion, THD , can be written as

$$THD = \frac{\sqrt{\sum_{k=2}^N |a_{kf}|^2}}{|a_f|} \quad (15)$$

3.3.2 Two-Tone Intermodulation

A two-tone intermodulation test uses two different frequencies to excite the system and determines non-linearities by calculating the complex amplitude/phase coefficients at the sum and the difference of these frequencies.

3.3.3 Transfer Matrix Measurements

Assume there is a linear relation between a set of channels \vec{a} and \vec{b} :

$$\vec{b} = M\vec{a} \quad (16)$$

with M the transfer matrix. Ideally, one would excite each a channel individually and measure the response b_n , thus, determining each row of M separately. In reality, the excitation might not be orthogonal in \vec{a} and one will end up with n measurements of both \vec{b}_n and \vec{a}_n . Choosing n different sine wave frequencies—one for each excitation—one can make all n measurement simultaneously by determining the corresponding transfer coefficients. If we form matrices from the measure vectors \vec{a}_n and \vec{b}_n ,

$$A_{nk} \equiv a_{nk} \quad \text{and} \quad B_{nl} \equiv b_{nl} \quad (17)$$

where the indices k and l are denoting the a and b channel numbers, respectively. Assuming that there are as many a channels as there are excitations, one can determine M by

$$M^T = A^{-1}B \quad (18)$$

where T denotes the transposed matrix. To be able to invert the matrix A none of the excitation signals (as measured by the a channels) can be co-linear with the others.

3.4 POLE-ZERO CURVE FITTING

TBD

3.5 CORRELATION MEASUREMENTS

TBD

3.6 TIME MEASUREMENTS

TBD

3.7 OCTAVE BAND ANALYSIS

TBD

4 TEST ORGANIZATION

Fig. 8 shows the interactions between a diagnostics test and the application program interfaces (APIs) which communicate with the user (storage API), the network data server (Real-Time Data Distribution API) for obtaining channel data, the excitation engine (excitation API) and the test point manager (test point API).

4.1 INTERFACES

4.1.1 Storage API

All data—including test parameters, raw channel data, result arrays and plot settings—are managed by the diagnostics storage object. Any data object can be accessed by its name. Typically, the user will set test parameters prior of starting a test. The diagnostics test program will retrieve these parameters and setup the test. While the measurement is under way channel data is received by the real-time data distribution interface and automatically saved in the storage

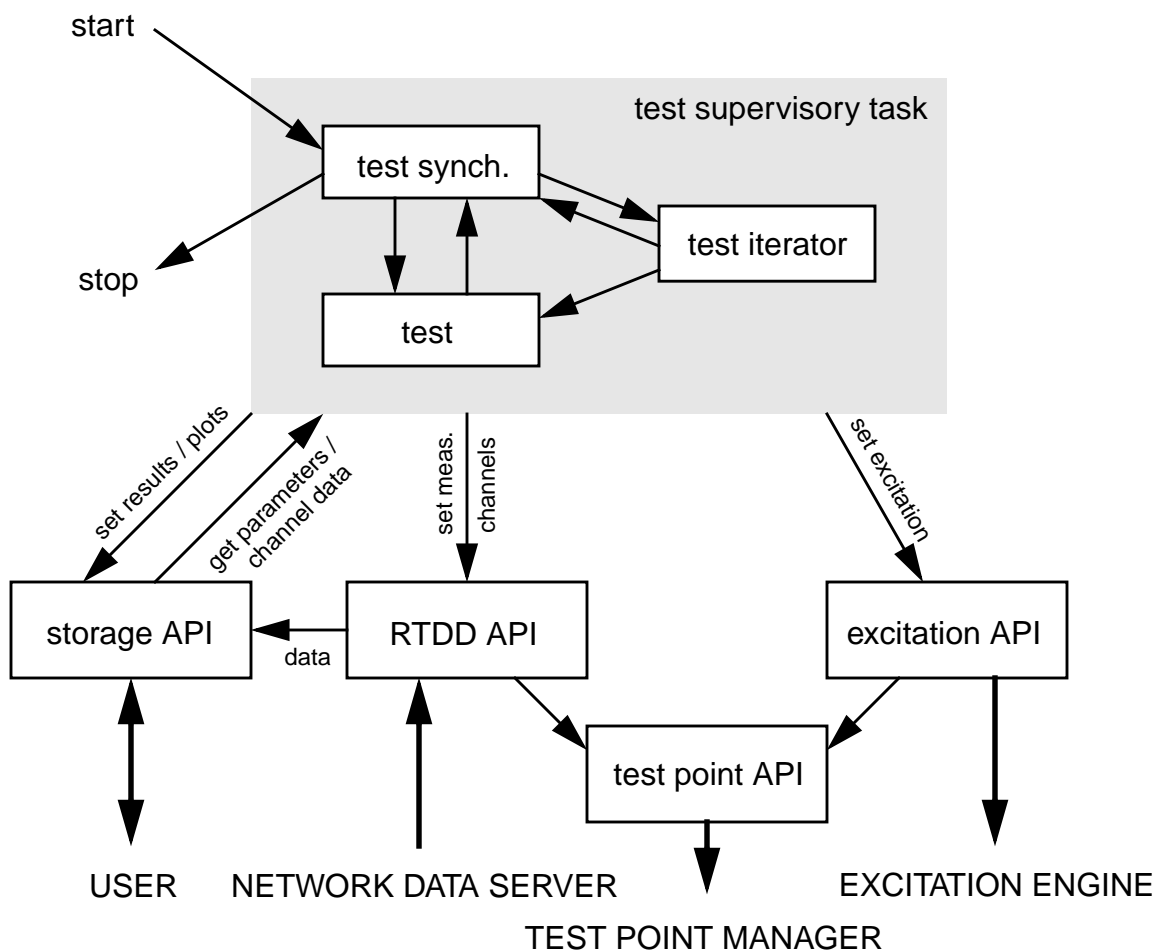


Figure 8: Interaction between a diagnostics test supervisory task and the interfaces to the real world.

object. After a measurement cycle has finished the analysis task will read the channel data, analyze it and write the results back.

4.1.2 Real-Time Data Distribution API

The real-time data distribution interface obtains data from the network data server on a continuous basis. It will subscribe the channels which are needed for the test analysis at the beginning of a test. While the test is running the data comes at a fixed rate of 16 Hz. The real-time data distribution interface has the capability to preprocess the data before sending it to the analysis algorithms (see Fig. 9). The first stage of the preprocessing is a time delay filter which can delay the data by an integer number of samples. This stage is followed by the first filter-decimation stage, the down-conversion stage and the second filter-decimation stage. If the down-conversion stage is included the data streams becomes complex. Before the data is partitioned into segments, the time series can be shifted forward in time to compensate for the delay accumulated in the decimation filters. The time shift has to be a multiple of the sampling period after decimation. Since the filter-decimation stages are based on FIR filters, the time delay accumulated within them is a multiple of the sampling period at their input, but not necessarily a multiple of the one at their output. To compensate for fractional time delays at the output stages, the delay filter at the input can be used to fine adjust the data stream. By default all filter delays are removed automatically and the partitioned data has no remaining time deviation.

4.1.3 Excitation API

The excitation API communicates with the excitation engine for applying excitation signals to the instrument. For a description of the excitation engine see Section 5.

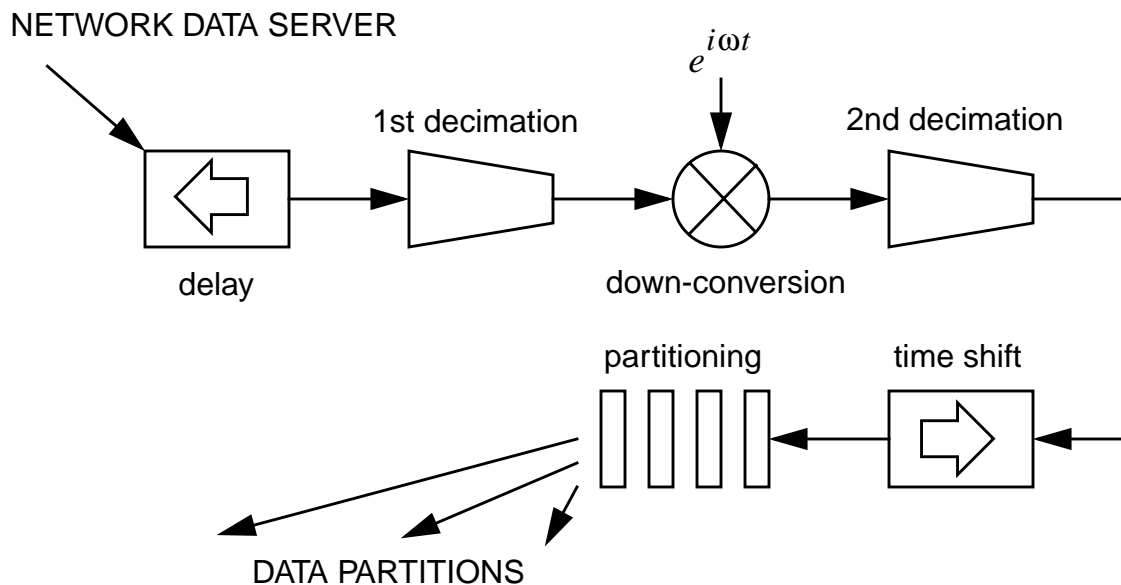


Figure 9: Data preprocessing in the real-time data distribution interface.

4.1.4 Test Point API

The test point API communicates with the test point manager which is able to activate digital test points in the data acquisition system. Digital test points are used to read auxiliary channels from the ISC digital servo systems and to inject digital waveforms. Test points are automatically selected by the RTDD API or the excitation API if the requested channel is accessible through a test point.

The test point API also enables and disables analog excitation inputs. (Most analog inputs have an enable/disable switch which is controlled through an EPICS channel.)

4.2 TEST SUPERVISORY

When a user starts a diagnostics test, the diagnostics kernel spawns a new test supervisory task. The supervisory task is selected by setting the variable “Supervisory” to the name of the supervisory task (see Appendix B.1.1). Currently, only one supervisory task with name “standard” is supported.

4.2.1 Standard Supervisory Task

In principle every test could be implemented as its own supervisory task. This has the disadvantage that parameter scans and optimizations would need an additional supervisory task for each test they can be combined with. Also, every test would have to implement its own synchronization means. The standard supervisory task provides a solution for this problem by separating the supervisory task into a synchronization task which is common to all tests, a test iterator and the test itself. Both the test iterator and the actual test are selectable by the user. This approach has the advantage that if any new test iterator is developed, it is immediately applicable to all tests; and if a new test is developed it can use all common test iterators. Fig. 10 presents the flow chart of the standard supervisory task.

The standard supervisory task implements a few basic synchronization tools:

- i)* the start time can be set using the variable “Sync.Start” (GPS nsec); default is immediate.
- ii)* a wait time can be specified with “Sync.Wait” (sec).
- iii)* the execution of a test or a test step can be triggered by an EPICS channels with “Sync.WaitForStart” and “Sync.WaitAtEachStep”, respectively.
- iv)* a test can send a trigger signal on an EPICS channels after it or one if its steps has finished; “Sync.SignalEnd” and “Sync.SignalEndOfStep”, respectively.

At initialization the supervisory task will setup the measurement environment as specified by the “Env[N]” variable sets (see also Appendix B.1.3).

A test can be aborted at any of the synchronization steps shown in Fig. 10. Currently, a the test can only be paused after an iterator evaluation has been completed.

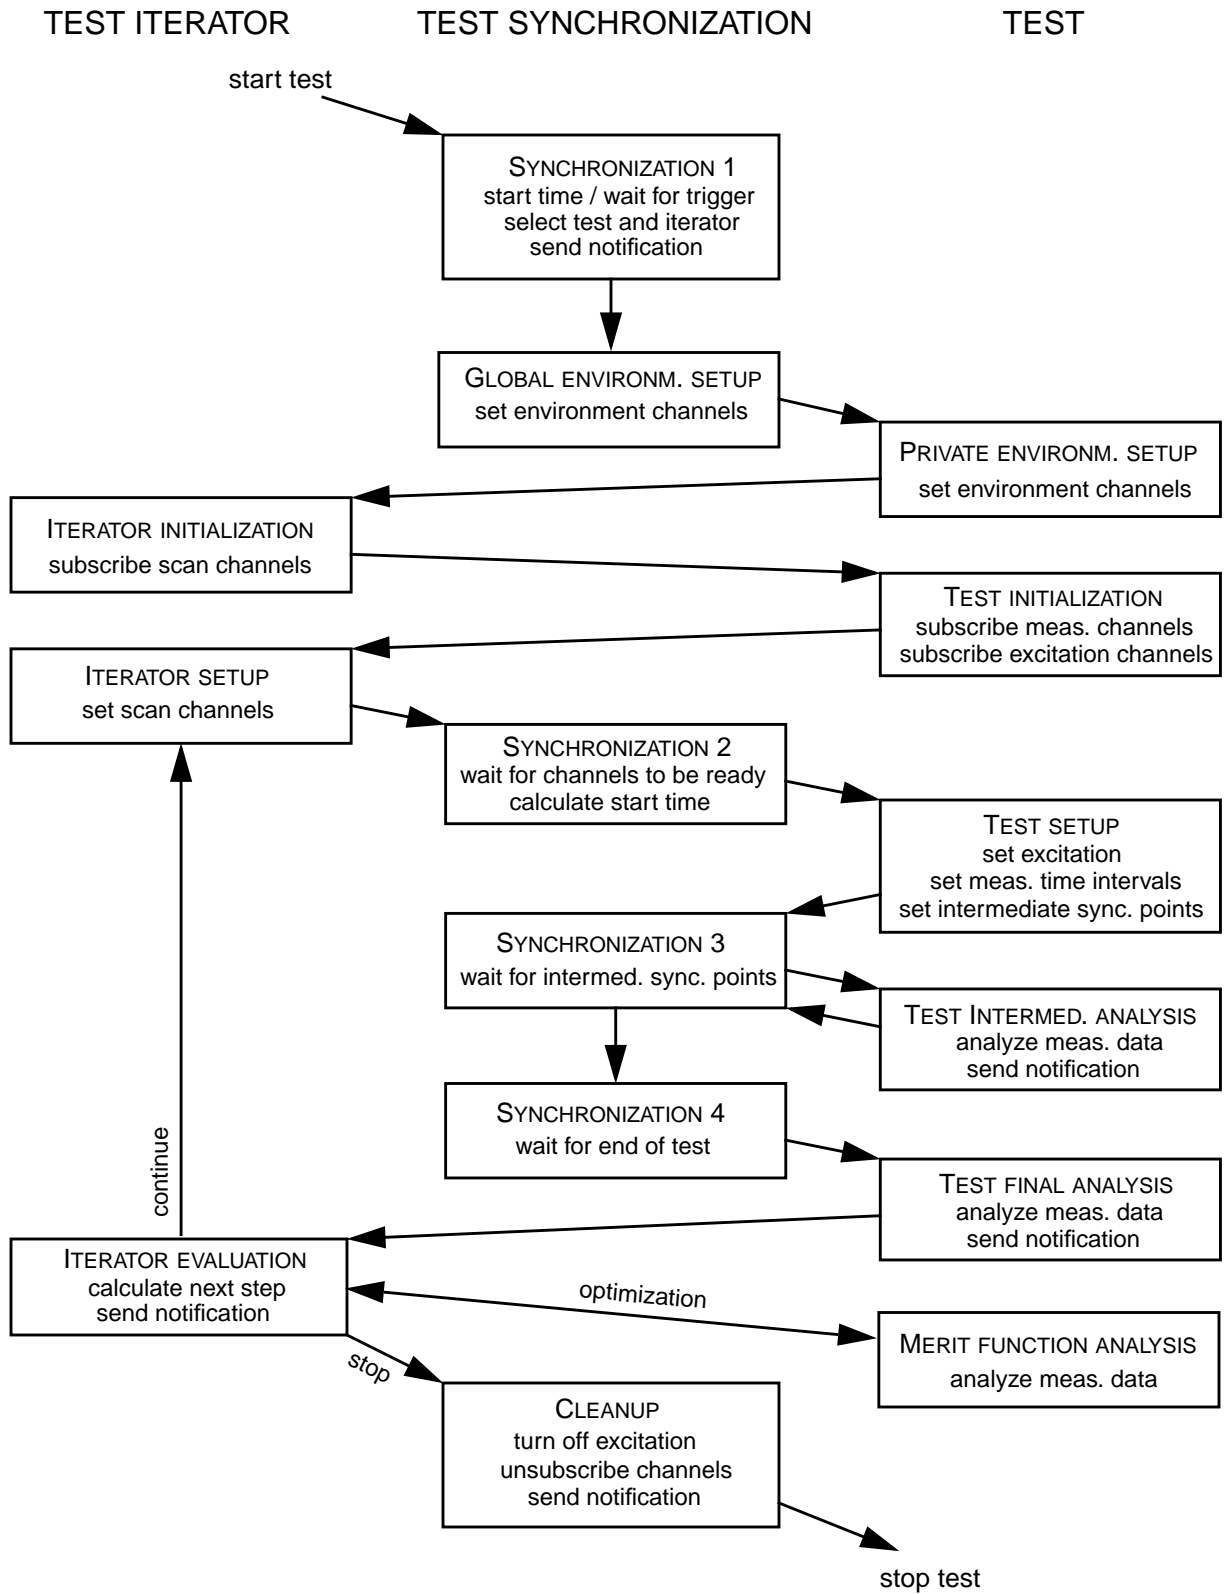


Figure 10: Flow chart of the standard diagnostics test supervisory task.

4.3 TEST ITERATORS

The purpose of test iterators are to repeat a standard test. In the simplest case the test is repeated with identical parameters. When performing a parameter scan the test is repeated while changing the value of a parameter. A parameter scan can be combined with an optimization to find the “best value” of a parameter set. The test iterator is selected by the variable “TestIterator” (see Appendix B.1.1).

4.3.1 Repeat

The default test iterator is the trivial one with a repetition value of one. The number of repeats can be set with the variable “Sync.Repeat”, whereas the repeat rate is determined by “Sync.RepeatRate”. This repeat rate is interpreted as a minimum time between tests. In other words the second test is not started before the first one has finished completely.

4.3.2 Parameter Scan

For a parameter scan at least one “Scan[N]” variable sets have to be defined and set active (see Appendix B.1.4). A test will be repeated while scanning a parameter of an excitation channel such as the amplitude, the offset or the frequency. An example would be to investigate the coupling of input beam jitter into the gravitational wave band (sine response test) while changing the angular alignment of the interferometer test masses.

4.3.3 Optimization

For an optimization to take place the “Find” variable set has to be defined and enabled. Additionally, a valid parameter scan has to be setup (see Appendix B.1.5). An optimization is a parameter scan where parameters are adjusted to their “best value”. Typically, the parameter is set to a new value after a scan interval has been completed. The “best value” is determined by a merit function which has to be provided by the test.

4.4 TESTS

4.4.1 Sine Response

The timing diagram of a sine response test is shown in Fig. 11. The full test time is divided into a dwell time and a number of measurement periods. The dwell time is divided into a ramp up period and a settling time (see also Appendix B.3.1). The settling time, t_S , is determined by the “Test.SettlingTime” values, t_Δ and n_c , according to the following formula (negative values are ignored):

$$t_S = \min\left(t_\Delta, \frac{n_c}{f}\right) \quad (19)$$

where f represents the frequency of the excitation signal. In case of multiple excitations with difference frequencies the smallest frequency value is used for the calculation.

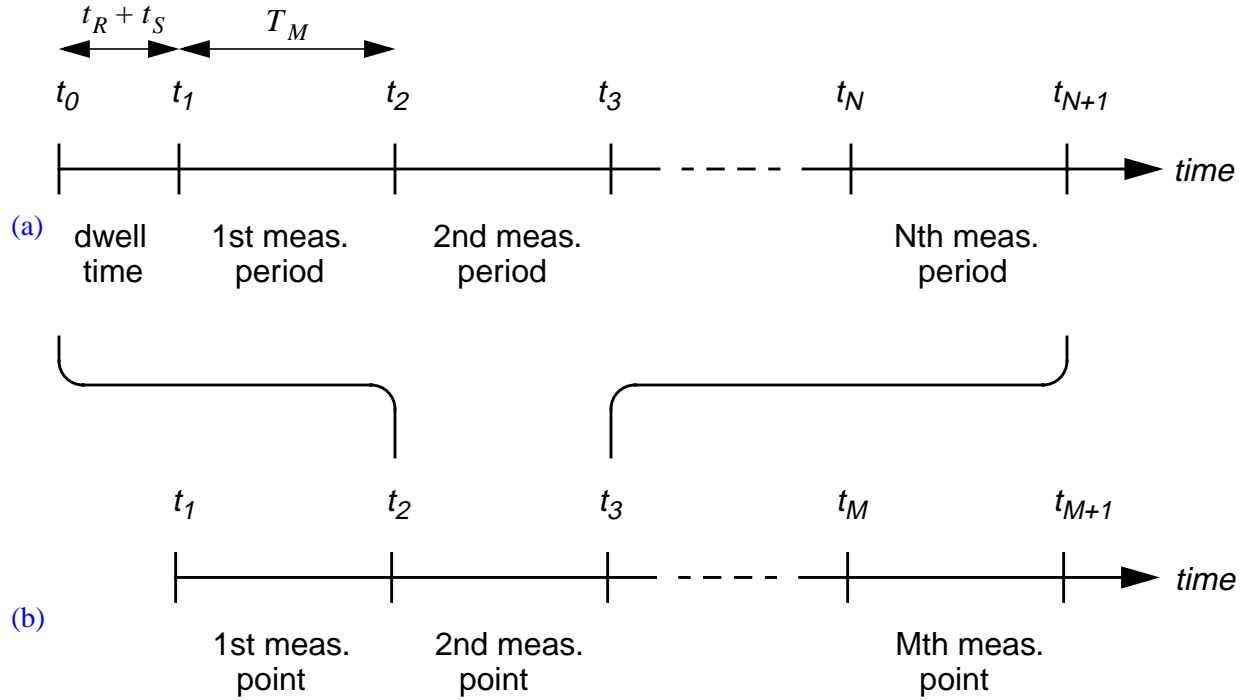


Figure 11: Timing diagram of a sine response (a) and a swept sine test (b). N and M are the number of averages and the number of measurement points, respectively.

During the ramp up time the amplitude of the excitation signals are slowly increase to their final requested levels. The ramp up time, t_R , is determined by the following formula:

$$t_R = \max(t_S, t_{\max}) \quad (20)$$

where $t_{\max} = 1$ sec is the maximum allowed ramp up time. The total dwell time $t_R + t_S$ is always rounded up to the next sampling period. If an excitation signal can not be synchronized with a GPS clock, the minimum ramp up time is set to 500ms to account for network latencies. The default ramp up function is a quadratic phase-in, see Eqns. (38) and (41).

The time of a single measurement T_M is determined by the “Test.MeasurementTime” values, T_Δ and N_c , according to the same formula as the settling time with the only modification that T_M is always rounded up the next full cycle and sampling period. If multiple excitation frequencies are present, the smallest frequency is used to perform the calculation.

4.4.2 Swept Sine

A swept sine is essentially a series of sine response tests with a single excitation frequency which is swept through a predefined frequency interval. The frequency points can be spaced linearly,

logarithmically or by user supplied values. The frequency points f_i are calculated from the start and stop frequencies, f_{start} and f_{stop} with the following formulae:

linear
$$f_i = f_{\text{start}} + (i - 1) \frac{f_{\text{stop}} - f_{\text{start}}}{M - 1} \tag{21}$$

logarithmic
$$f_i = f_{\text{start}} \left(\frac{f_{\text{stop}}}{f_{\text{start}}} \right)^{\frac{i-1}{M-1}} \tag{22}$$

where M represents the number of frequency points and the index i runs from 1 to M .

The dwell time is calculated the same way is for a simple sine response test. The only difference is that the ramp up signal between measurement points is replaced by a ramp signal which guarantees a smooth transition in both amplitude and frequency. For the amplitude Eqns. (38) and (41) are used again; for sweeping the frequency Eqns. (42) and (43) are used for linear swept sine tests and Eqns. (44) and (45) for logarithmic swept sine tests, respectively.

4.4.3 FFT Tests

The timing diagram of an FFT test is shown in Fig. 12. The total test time is broken into possibly overlapping measurement intervals. The frequency span, f_{span} of the FFT is calculated by

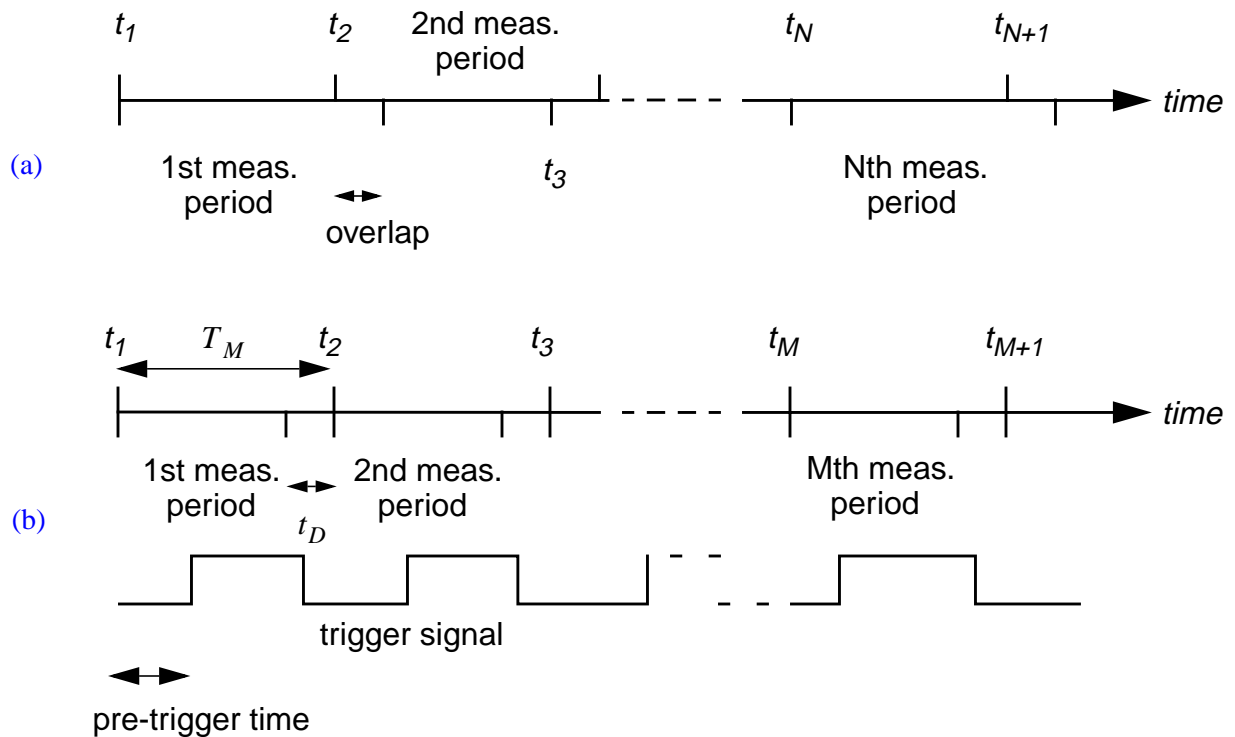


Figure 12: Timing diagram of an FFT test (a) and a time series/trigger response test (b). N and M are the number of averages and the number of measurement points, respectively.

rounding $S_{\text{stretch}}(f_{\text{stop}} - f_{\text{start}})$ to the next frequency which is a power of 2. The stretch factor, S_{stretch} , is set at 1.1 which guarantees that the requested frequency range is within the bandwidth of the decimation filters. To make sure that the number of FFT points is a power of 2, the specified bandwidth is rounded to the closest power of 2. The measurement time is determined by the bandwidth through

$$T_M = \frac{1}{f_{BW}} \quad (23)$$

Of course, if a non-uniform window function is selected, the effective measurement bandwidth is given by the width of the window function, i.e.

$$f_{BW}^{\text{eff}} \approx \frac{1}{\Delta t_{\text{window}}} \quad (24)$$

If $f_{\text{start}} > 0$, the time series is first down-converted by

$$f_{\text{zoom}} = f_{\text{start}} + N_{FFT} \text{round}\left(\frac{f_{\text{stop}} - f_{\text{start}}}{2N_{FFT}}\right) \quad (25)$$

where N_{FFT} is the number of points used by the FFT algorithm. It is given by

$$N_{FFT} = \begin{cases} 2 f_{\text{span}} / f_{BW} & \text{for } f_{\text{start}} = 0 \\ f_{\text{span}} / f_{BW} & \text{for } f_{\text{start}} > 0 \end{cases} \quad (26)$$

In general, input channels pass through three stages of preprocessing. First, a decimation stage adjusts the sampling rate to the required signal bandwidth. Second, the time series can be down-converted by f_{zoom} ; and third, an additional decimation stage adjusts the sampling rate to the frequency span of the FFT.

The following rules apply for determining sampling rates and decimation factors:

(i) The sampling rate, f_{sample} , after the first stage is deduced from $\max(2f_{\text{stop}}, 2f_{\text{span}})$ by rounding to the next power of 2.

(ii) The decimation factor of the second stage is set to 1 if no down-conversion was applied and to $f_{\text{sample}} / f_{\text{span}}$ otherwise.

Compared to a standard (real) FFT analysis, a zoom analysis works with a complex time series which effectively doubles the number of data values, but only needs half the sampling rate (since negative frequency components are truly different from positive ones). As a net effect the number of data values sent to the FFT algorithm are identical.

If the number of averages is greater than 1, the start time t_i for each measurement interval is given by

$$t_i = t_o + i(1 - \eta)T_M \quad (27)$$

with η the overlap factor.

4.4.4 Time Series Measurements

Time series and trigger response measurements are broken up in equal length measurement intervals (see Fig. 12). The trigger rate, f_{trigger} , is related to the measurement time, T_M , by

$$f_{\text{trigger}} = \frac{1}{T_M} \quad (28)$$

The stored time series can be limited in length by specifying a dead time, t_D , which is subtracted from the end of the measurement interval. In order to keep the averaging algorithms simple both the length of the measurement interval and the settling time have to be a multiple of the sampling period after decimation. If the Nyquist frequency of the original time trace is more than twice as large as the required bandwidth, the data is filtered and decimated to a rate no less than twice the bandwidth and no more than four times the bandwidth. The pre-trigger time specifies the time the start of the trigger signals lags behind the start of the measurement.

5 EXCITATION ENGINE

5.1 OUTPUT WAVEFORMS

5.1.1 Overview

Every output signal produced by the arbitrary waveform generator consists of two parts: a fundamental waveform and phase-in and phase-out transition. The following fundamental waveforms are supported by the signal generator:

$$\text{sine wave} \quad s(t) = A \sin(2\pi ft) \quad (29)$$

$$\text{square wave} \quad s(t) = \begin{cases} +A & 0 \leq 2\pi ft \bmod 2\pi < \pi \\ -A & \pi \leq 2\pi ft \bmod 2\pi < 2\pi \end{cases} \quad (30)$$

$$\text{ramp} \quad s(t) = A \frac{\phi}{2\pi} \text{ with } \phi = 2\pi ft \bmod 2\pi \quad (31)$$

$$\text{triangle} \quad s(t) = \begin{cases} A \left(\frac{2\phi}{\pi} - 1 \right) & 0 \leq 2\pi ft \bmod 2\pi < \pi \\ A \left(3 - \frac{2\phi}{\pi} \right) & \pi \leq 2\pi ft \bmod 2\pi < 2\pi \end{cases} \quad (32)$$

$$\text{flat noise} \quad (33)$$

$$\text{band limited noise} \quad (34)$$

$$\text{pink noise} \quad (35)$$

$$\text{sweep} \quad (36)$$

5.1.2 Periodic Waveforms

The output signal $o(t)$ of an arbitrary waveform generator for a periodic waveform starting at time t_0 , lasting for a duration Δt , ramping up between t_0 and $t_0 + t_{PI}$, and phasing out after $t_0 + t_2$ ($t_2 = \Delta t - t_{PO}$) can be written as:

$$o(t) = s(2\pi f(t - t_0) - \phi_0 + \varphi(t - t_0 - t_2, t_{PO})) \times b(t - t_0, \Delta t, t_{PI}, t_{PO}) \quad (37)$$

where $s(\phi)$ is the periodic signal (fundamental waveform) of frequency f and phase shift ϕ_0 . The functions φ and b are then functions which determine the phase-in and phase-out transitions at

the beginning and at the end of the output signal. The phase-in and phase-out times are denoted with t_{PI} and t_{PO} , respectively. The function b controls the amplitude and we write it as:

$$b(t-t_0, \Delta t, t_{PI}, t_{PO}) = \begin{cases} 0 & t < 0 \text{ or } t > \Delta t \\ 1 & t_{PI} \leq t < \Delta t - t_{PI} \\ g_1(t, t_{PI}) & 0 \leq t < t_{PI} \\ g_2(t-t_2, t_{PO}) & \Delta t - t_{PO} \leq t \leq \Delta t \end{cases} \quad (38)$$

where g_1 and g_2 are the amplitude phase-in and phase-out functions, respectively. The following options are available: step, linear ramp, and quadratic.

$$\begin{array}{l} \text{step} \\ \end{array} \quad \begin{array}{l} g_1(t, t_{PI}) = 0 \\ g_2(t, t_{PO}) = 1 \end{array} \quad (39)$$

$$\begin{array}{l} \text{linear} \\ \end{array} \quad \begin{array}{l} g_1(t, t_{PI}) = \frac{t}{t_{PI}} \\ g_2(t, t_{PO}) = 1 - (1-c)g_1(t, t_{PO}) \end{array} \quad (40)$$

$$\begin{array}{l} \text{quadratic} \\ \end{array} \quad \begin{array}{l} g_1(t, t_{PI}) = -\left(\frac{t}{t_{PI}}\right)^4 + 2\left(\frac{t}{t_{PI}}\right)^2 \\ g_2(t, t_{PO}) = 1 - (1-c)g_1(t, t_{PO}) \end{array} \quad (41)$$

where c denotes the amplitude ratio of the signal before the phase-out and the one after. For two periodic waves of different frequencies which should smoothly transform into each other, it is usually more practical to sweep the frequency, rather than phase-out one signal by ramping it down while ramping up the other one. The sweep function φ is defined as follows:

$$\varphi(t, t_{PO}) = g_1(t, t_{PO})(2\pi\Delta f t - \overline{\Delta\phi}) \quad (42)$$

with $\Delta f = f_2 - f_1$ the frequency difference between the beginning and the end of the sweep. The phase adjustment $\overline{\Delta\phi}$ can be written as:

$$\overline{\Delta\phi} = \Delta\phi + 2\pi f \Delta t + 2\pi \Delta f t_{PO} \Big|_{\text{mod } 2\pi} \quad (43)$$

with $\Delta\phi$ the phase difference between the signal at $t = t_0$ and $t = t_0 + \Delta t$.

The above sweep function can be readily generalized to logarithmic sweeps. For the time during the sweep one can rewrite Eqn. (37) as follows:

$$o(t) \Big|_{t_2 \leq t - t_0 \leq \Delta t} = s(2\pi f t_2 - \phi_0 + \varphi_{\log}(t - t_0 - t_2, t_{PO})) \quad (44)$$

with the sweep function

$$p_{\log}(t, t_{PO}) = \left[2\pi f t_{PO} \left(\frac{f + \Delta f}{f} \right)^{\frac{t}{t_{PO}}} - \overline{\Delta\phi} \right] \frac{t}{t_{PC}} \quad (45)$$

where the phase adjustment is the same as in Eqn. (43).

5.1.3 Non-Periodic Waveforms

Non-periodic output waveforms can be written as:

$$o(t) = n(t - t_0) \times b(t - t_0, \Delta t, t_{PI}, t_{PO}) \quad (46)$$

with the only difference to a periodic waveform that the actual signal $n(t)$ does depend on the time directly rather than the phase, and that no sweep function is provided.

APPENDIX A NETWORK INTERFACES

A.1 SERVICES INFORMATION

Diagnostics services such as the test point interface, the excitation engine and the diagnostics test kernel implement a configuration and information interface which can be queried by a client using broadcast messages on the local subnet. This allows user interface programs to be configured dynamically. Service information requests and answers all have the same format:

Service information request/answer:

msb			lsb	1	2		N
-----	--	--	-----	---	---	--	---

information id

ASCII string (optional)

The transport protocol is udp/ip. A service information server can answer with any number of datagrams to a single request.

The following information requests are available:

info id	argument	description	answer	answered by
0	-	ask for interfaces	interface descriptor	all

An interface descriptor has the following format:

'interface' 'ifo' '#' 'host name' 'port/prog #' 'version'

Examples are:

```
awg 0 0 gdsawg1 8220876582 1
awg * 0 cobox0 5000 *
```

The excitation engine and the test point manager implement their own information server for their respective services, whereas the diagnostics kernel relies on daemon to announce its services. The following lines should be added to the '/etc/inet/services' and '/etc/inet/inetd.conf' files on every machine which runs a remote diagnostics kernel:

```
diagconf 5355/udp
diagconf dgram udp wait gds /home/gds/bin/diagconfd diagconfd \
/home/gds/param/diag.conf
```

(Assuming the full path name of the diagnostics service daemon is '/home/gds/bin/diagconfd' and the corresponding configuration file is located at '/home/gds/param/diag.conf'.)

Each line of the configuration file contains an interface descriptors. Lines starting with '#' are ignored. Service names starting with a '?' will only be part of the configuration information, if the corresponding machine is alive (i.e. answers to a ping). Service names starting with '&' are using

host name lookup to replace hostnames with IP names. A service name starting with ‘&?’ will first look up the host name and then wait for the ping answer.

A.2 MESSAGE PASSING INTERFACE

The message passing interface is the main communication channel between the user interface and the diagnostics test kernel.

A.2.1 USING REMOTE PROCEDURE CALLS

The rpc message passing interface implements four basic functions:

- i)* open a connection to a diagnostics kernel,
- ii)* close the connection to a diagnostics kernel,
- iii)* send a message to a diagnostics kernel and wait for the answer, and
- iv)* install a callback function for messages send back by the diagnostics kernel.

These functions are defined in the C header file ‘gdsmsg.h’ (for a more detailed description see there, or in the corresponding web page). A message consists of a message header and a parameter argument; it returns a reply argument. Valid message headers are ASCII strings which are generally identical to the commands which are recognized by the diagnostics kernel as specified in section 2.3. In general, a command and its ASCII arguments are passed through the message header, whereas binary arguments are passed through the parameter argument.

One important difference to the command line interface are the ‘save’ and ‘restore’ commands: *(i)* the flag argument must be specified, *(ii)* the filename is assumed to be local to the machine which runs the diagnostics kernel, and *(iii)* the reserved filename specifier ‘-’ is used to indicate that file is part of the message. When using ‘save’, the file is passed back through the reply argument. When using ‘restore’ the file has to be passed to the diagnostics kernel through the parameter argument.

The diagnostics kernel can run on a local machine—in this case the shared object library ‘libgds.so.1’ is loaded dynamically when the connection is established—or it can run on a remote machine—in this case the diagnostics kernel has to run on the remote machine. The diagnostics server listens at a well known tcp/ip port for a client to connect. Every client launches a new diagnostics kernel which then established an rpc communication channel. Upon on connection the client sends three 32 bit number in network byte order:

1. If the first number is larger than 65535, it is interpreted as an rpc program number which will be used to send notification messages back to the client,
2. The second number is a client specific identification number. It will be used when sending notification messages, and
3. The third argument is a flag describing the required services (as specified by the gdsMsgOpen function).

In return the server will answer with up to three 32 bit numbers in network byte order:

1. The first one is the status flag which describes whether the server was initialized successfully. If non-zero, it indicates an error and the transmission is closed,
2. When successful, the second number describes the rpc program number which has to be used by the client for subsequent calls, and

3. The third argument describes the capabilities provided by the server (as defined by the `gdsMsgOpen` function).

The diagnostics server program can be started either at the command line prompt or through a port monitor. Assuming the current path is the diagnostics binary directory the server is started by typing `./diagd`. When using the port monitor, the following lines have to be added to the `/etc/inet/services` and `/etc/inet/inets.conf` files:

```
diag 5354/tcp
diag stream tcp nowait gds /home/gds/bin/diagd diagd
```

(Assuming the full path name of the diagnostics server is `/home/gds/bin/diagd`.)

Whenever a request is made at the specified port a new diagnostics kernel is launched and a private two way communication channel is established between the caller and the diagnostics kernel. After closing the connection the diagnostics kernel is automatically terminated. To prevent an automatic shutdown, the client has to call the keep alive function every 30 sec.

A.2.2 USING SOCKETS

Instead of using remote procedure calls a client can also request that messages are sent through tcp/ip sockets. The start-up procedure is exactly the same as for the rpc interface, but with a value for the first argument between 0 and 65535. If a non-zero number is specified, it is interpreted as a port number of a socket maintained by the client, and it will be used by the server to send back notification messages. A zero value indicates that the server should use the already open connection to both transfer and receive messages. The server will answer with a status number and, if successful, the flag describing the available services.


After establishing a connection to the diagnostics server messages are passed using the format shown below:

Header:




size of message header ASCII string (diagnostics command)

Binary argument:



size of binary argument binary argument

Return argument:



status size of binary argument ASCII / binary

Every message consists of a header and a binary argument (their sizes can be zero); it will return either an ASCII or a binary argument dependent on the command. Each argument is preceded by a 32 bit number in network byte order describing the length of the following argument. The server will send the answers back in the order the commands were received. The return argument consists of a status word, a word describing the following data and the data itself. A status of 0 marks a success, a status of -1 indicates an error, and a status of -2 indicates a notification message. (Notification messages are sent from the diagnostics kernel to the user interface using the same format but without at return argument.)

To prevent automatic shutdown of the server the client has to send at least one message every 30 seconds. If no regular message is available the client can send an empty message with header and argument of size zero (a null message will not produce any answer). Starting a diagnostics kernel for socket use is identical to the one for remote procedure calls, both from the command line and through a port monitor. There is one important difference between using rpc or sockets: the rpc interface is fully multi-threaded, i.e. several messages can be sent to diagnostics kernel by the same client independently, whereas messages sent through sockets are processed one after the other.

APPENDIX B DATA REPRESENTATION

The diagnostics kernel manages variables, parameters and data which describe a diagnostics test using the same basic representation. The basic storage objects are data objects and parameter objects. Parameter objects consists of a name, a type and a value or a list of values. A data object is a container which can contain a multidimensional data array, a set of parameter objects and other data objects. The diagnostics kernel stores all variables associated with a diagnostics test within a single (global) data object. The diagnostics kernel limits the hierarchical levels of data objects to two. Meaning, the global object can contain parameters and data objects which in turn can contain other parameter objects, but no more data objects.

Every data object and every parameter object within a single data object must have a unique name. Any parameter can then be accessed by specifying its name and the name of its data object (group name). The parameters in the global data objects can be accessed by their name only, or by specifying an empty data object name. A name is not case sensitive and can not contain the ‘.’ character, since this character is used to separate its name from the group name.

A name must not contain square brackets, because they are used for indexing. Names and group names can have an array index, e.g., ‘Scan[1]’ or ‘H0:GDS_TEST[0][2]’. A maximum of two array indices are supported with values ranging from 0 to 999.

Both data and parameter objects have associated data fields consisting of a data type, a dimension list, the actual datum, an optional string describing the physical unit and an optional comment string. Supported data types are:

Name	C/C++ data type	Abbreviation
8 bit integer	char	c
16 bit integer	short	s
32 bit integer	int	i
64 bit integer	long long	ll
boolean	bool	b
single precision floating point	float	f
double precision floating point	double	d
single precision complex number	complex<float>	zf
double precision complex number	complex<double>	zd
string	char*	st
channel name	char*	ch

Parameter objects typically have a single value associated with them, whereas data objects can be multi-dimensional.

The diagnostics kernel uses a set of data objects and parameters with predefined names for describing: (i) the test parameters, (ii) the test results, (iii) the raw data and (iv) the plot options. The following data objects are used by the diagnostics kernel:

Name	Type	Description
Def	test param.	Global default settings of diagnostics tests.
Sync	test param.	Synchronization information.
Env[N]	test param.	Excitation environment; N from 0 to 99.
Scan[N]	test param.	Scan parameters; N from 0 to 9.
Find	test param.	Optimization parameters.
Test	test param.	Specific settings for a diagnostics test.
Index	result	Index of the results
'channel name'[M][N]	raw data	measured time series; M and N from 0 to 999.
Result[N]	result	results of diagnostics test: FFT, transfer functions, averaged time series, coefficients and list of measurement values; M from 0 to 999.
Plot[N]	plot settings	result plots; M from 0 to 99.

All data objects have a string parameter of name `ObjectType` which identifies the data object. The following object types are supported:

ObjectType	Associated data objects
DiagnosticsTest	global
Defaults	Def
Synchronization	Sync
Environment	Env[N]
Scan	Scan[N]
Optimization	Find
TestParameter	Test
TimeSeries	'channel name'[M][N] or Result[N]
Spectrum	Result[N]
TransferFunction	Result[N]
Coefficients	Result[N]
MeasurementTable	Result[N]
Plot	Plot[N]

B.1 COMMON

This section list the parameters which are common to all diagnostics tests.

B.1.1 GLOBALS

A few parameters are defined in global scope:

Name	Type	Dim	Description
ObjectType	st	1	DiagnosticsTest
TestType	st	1	describes the diagnostics test class. Possible values are: TimeSeries, SweptSine, FFT, SineResponse, RandomResponse.
TestName	st	1	name of test, user supplied.
Supervisory	st	1	name of supervisory task; default standard.
TestIterator	st	1	name of test iterator; default repeat. Possible values are: repeat, scan, find.
Comment	st	1	user supplied comment.
TestTime	ll	1	time when test was done in GPS nsec.
TestTimeUTC	st	1	time when test was done in UTC format, e.g. 1998-11-08 17:40:00.032035.

The following default parameters are common to all tests:

Name	Type	Dim	Description
Def.ObjectType	st	1	Defaults
Def.AllowCancel	b	1	if true (default), cancel a test is allowed at any time.
Def.NoStimulus	b	1	if true, no stimulus is applied; default is false.
Def.NoAnalysis	b	1	if true, omits the analysis; default is false.
Def.KeepTraces	i	1	specifies how many of the original time traces are kept (-1 represents all); default 100.
Def.SiteDefault	c	1	describes the default site; any channel name that contains an 'X' as the site identifier will be adjusted to the specified default site.
Def.SiteForce	c	1	overrides the site identifier in every channel name.
Def.IfoDefault	c	1	describes the default interferometer number. Any channel name containing an 'X' instead of an interferometer number will be adjusted to the specified interferometer.
Def.ifoForce	c	1	overrides any none zero interferometer identifier in every channel name.

Channel names can be specified site and interferometer independent. To do so the character ‘X’ is used instead if the site or the interferometer identifier, respectively. Any ‘X’ character will be replaced with its corresponding default value before the test is started. Alternatively, it is possible to override site or interferometer identifiers of all channel names. This can be useful when moving a diagnostics test from one instrument to an other.

B.1.2 SYNCHRONIZATION TOOLS

Diagnostics test can be synchronized with EPICS channels. A test can halt at the beginning of a test and/or at each step of a measurement until a trigger signal is received on a specified EPICS channel. Similarly, a diagnostics test can send a trigger signal on a specified EPICS channel at the end of a test and at the end of each measurement step.

Name	Type	Dim	Description
Sync.ObjectType	st	1	Synchronization
Sync.Start	ll	1	start time of test (GPS nsec); default 0 (now)
Sync.Wait	d	1	time to wait before starting test (in sec); default 0. the wait time is added to the start time.
Sync.Repeat	i	1	number of times of repeating the test; default 1.
Sync.RepeatRate	d	1	rate of repeating tests (in sec); default 0 (no wait).
Sync.WaitForStart	ch	1	if defined, the test only starts after receiving a trigger signal on the specified EPICS channel (waits for a 1; resets channel to zero after trigger was received).
Sync.WaitAtEachStep	ch	1	if defined, the test only starts a new step after receiving a trigger signal on the specified EPICS channel.
Sync.SignalEndOfStep	ch	1	if defined, signals the end of a test step on the specified EPICS channel (sets channel to 1).
Sync.SignalEnd	ch	1	if defined, signals the end of the test on the specified EPICS channel.

B.1.3 ENVIRONMENT

While performing a diagnostics test, excitation and EPICS channels can be set to specified waveforms and values, respectively. This environment is set before the test starts and automatically reset after the test terminates. Environments are numbered, describing a single excitation channel each:

Name	Type	Dim	Description
Env[N].ObjectType	st	1	Environment
Env[N].Active	b	1	if true (default), includes the Env[N] excitation channel.
Env[N].Channel	ch	1	string describing an excitation channel name.

Name	Type	Dim	Description
Env[N].Waveform	st	1	string describing an excitation waveform. The format is the same format as described in Section 2.3.2.
Env[N].Points	f	M	describes the values of an arbitrary waveform.
Env[N].Wait	d	1	time to wait for the environment to settle down.

If an excitation channel is defined in the environment and is also used by a stimulus response test, the two excitation waveforms are simply added. Certain restrictions apply if the excitation channel is a stand-alone signal generator.

B.1.4 PARAMETER SCAN

Any test can be repeated while sweeping either frequency, amplitude or offset of one or multiple excitation channels. Sweep parameters are numbered from $N = 0, 2, \dots, 9$, describing a single excitation channel each:

Name	Type	Dim	Description
Scan[N].ObjectType	st	1	Scan
Scan[N].Active	b	1	if true (default), includes Scan[N] in the sweep.
Scan[N].Channel	ch	1	name(s) of sweep channel, can be an excitation channel, or an EPICS analog output.
Scan[N].Type	i	1	0 – linear sweep (default), 1 – logarithmic sweep, 2 – user supplied sweep points.
Scan[N].Direction	i	1	0 – upwards (default), 1 – downwards.
Scan[N].Parameter	i	1	0 – offset (default), 1 – amplitude, 2 – frequency Only offset is valid for an EPICS channel.
Scan[N].Frequency	d	1	specifies the frequency of the signal (not used, if the frequency is swept, or if EPICS channel).
Scan[N].Amplitude	d	1	specifies the amplitude of the signal (not used, if the amplitude is swept, or if EPICS channel).
Scan[N].Offset	d	1	specifies the offset of the signal (not used, if the offset is swept, or if EPICS channel).
Scan[N].Start	d	1	specifies the start point of the sweep (not used when user supplies sweep points).
Scan[N].Stop	d	1	specifies the stop point of the sweep (not used when user supplies sweep points).
Scan[N].N	i	1	specifies the number of points of the sweep (not used when user supplies sweep points).
Scan[N].Points	f	M	specifies the user supplied sweep points
Scan[N].Wait	d	2	specifies the settling time before each measurement; default is 10 cycles, or 65 ms, whatever is shorter. first value is time in sec.

The dimension of the scan is determined by the number of defined channel names. If fewer than N values are specified for an N dimensional parameter, the remaining values are automatically padded with zero.

A multi-dimensional linear or logarithmic sweep will use a multi-dimensional mesh to set the sweep points. For a user defined multi-dimensional sweep the user has to supply every sweep point which will be used; they can be located anywhere in the multi-dimensional parameter space.

If the same excitation channel which is defined in the environment and/or by a stimulus response test, is also used by a sweep, the excitation waveforms are simply added. Certain restrictions apply if the excitation channel is a stand-alone signal generator

B.1.5 PARAMETER OPTIMIZATION

A parameter scan can be combined with a parameter optimization. When optimization is enabled the sweep will try to set the sweep channel values to the best values as determined by a merit function. For a multi-dimension linear or logarithmic sweep the optimization process will scan through each dimension separately, and set the signal to the best value after each scan. A multi-dimensional sweep with user supplied sweep points is interpreted as a one dimensional optimization along the path of the user supplied points.

Name	Type	Dim	Description
Find.ObjectType	st	1	Optimization
Find.Enable	b	1	if true, enables the optimization process; must be selected together with a sweep. Default is false.
Find.Change	b	1	if true (default), the detector state at the end of the test will be set back to the values before the test, i.e. the optimization process tries to find an optimal parameter set without changing to the new and better state.
Find.Type	i	1	0 – maximum, 1 – minimum, 2 – zero, 3 – value
Find.Value	d	1	value to find, if type is 'value'.
Find.Function	i	1	The merit function which is used depends on the selected diagnostics test: SineResponse (single frequency only): 0 – amplitude, 1 – phase, 2 – harmonic distortion, 3 – intermodulation product (added in quadrature) TimeSeries: 0 – average, 1 – rms FFT or RandomResponse: 0 – band-limited power
Find.Param	d	2	low and high frequency values for band-limited power
Find.Method	i	1	0 – scan

B.2 RESULTS

A diagnostics test can consist of multiple measurements such as the individual steps of a sweep. Each measurement can contain multiple measurement points, such as the frequency points of a swept sine measurement. Typically, there is the need to store results of individual measurement points, as well as the results of each measurement step and the final result. This is achieved by using array indices which are part of the result name. However, there is no one-to-one relation between result index and measurement step or point. Instead, an index is used to list measurement results together with the information where they are stored.

B.2.1 INDEX

An index is automatically generated by the diagnostics kernel while performing a test. It is updated whenever a new results becomes available. The index is structured as follows:

Name	Type	Dim	Description
ObjectType	st	1	Result
Entry[N]	st	1	index entry
Status	st	1	current status describing the current measurement point/average and the current test iteration step.

An index entry has the following format: “<category>[step] : {<parameter> = <value> ;}” where the curly brackets denotes items which can be repeated multiple times. Possible categories are:

Index Category	Description
MasterIndex	list of all index entries (master index)
TimeSeries	describes a time series result
PowerSpectrum	describes a power spectrum measurement
Coherence	describes a coherence estimate for FFTs
CrossCorrelation	describes a cross-spectrum measurement
TransferFunction	describes a transfer function
CoherenceFunction	describes the coherence estimate for swept sine measurements
TransferCoefficients	describes a set of transfer coefficients
TransferMatrix	describes a transfer matrix
CoherenceCoefficients	describes a set of coherence coefficients
HarmonicCoefficients	describes a set of harmonic coefficients
IntermodulationCoefficients	describes a set of two-tone intermodulation coefficients

B.2.2 PLOT

A plot object does not contain the data it displays, rather it describes the settings of the plot and stores pointers to the data objects which contain the plotted data. Plot objects are not generated by

the diagnostics test, but rather by the graphical user interface to keep track of the plot settings. A plot object can handle A and B channels. When displaying a transfer function or a cross-power spectrum, A and B channels point to the individual measurement channel; and the plotting routine is responsible to calculate the ratio or product, respectively. The following options are used to control the appearance of the plot.

Parameter Name	Type	Dim	Description
ObjectType	st	1	Plot
PlotType	i	1	0 – XY plot (line/scatter), 1 – time series (line/scatter), 2 – power spectral density (PSD), 3 – Bode plot (dB/phase plots, line), 4 – histogram (lego). default: 0.
Trace	i	16	0 – none 1 – A channel 2 – B channel 3 – cross spectrum (PSD) 4 – transfer function B/A (PSD, Bode) 5 – coherence (PSD, Bode) 6 – normalized variance (PSD, Bode) default: 1.
PlotLabel	st	1	Plot label. Use TEX convention for greek letters, superscript and subscript, e.g. '\$\alpha_1\$' default: empty string (no label).
AspectRatio	d	1	aspect ratio of plot. default: 0.7.
Frame	b	1	if true, plot is framed. default: true.
Label[N]	st	1	Frame labels, numbering: clockwise starting at lower x-axis. default: empty strings (no labels).
Font	i	1	0 – Helvetica, 1 – Times. default: 0.
FontSize	d	3	Font sizes of plot label, frame labels and axes label, respectively (in relative units to the diagonal plot length). default: 0.045, 0.04, 0.03.
Axes	b	2	whether to draw x-axis (first value) and y-axis (second value). default: true, true.
AxesType	i	2	0 – linear, 1 – log. default: 0, 0.

Parameter Name	Type	Dim	Description
AxesPlacement	d	1	place of axes (in relative units to the diagonal plot length); positive unit – axes are placed outside default: 0.05.
AxesThickness	d	1	thickness of axes (in relative units to the diagonal plot length). default: 0.004.
AxesLabel	i	2	0 – no axes labels, 1 – labels at major ticks, 2 – labels at major and minor ticks. default: 1, 1.
MajorTicks	i	2	0 – automatic, 1 – no major ticks, 2 – user specified. default: 0.
MajorTicksStyle	i	2	0 – both directions, 1 – inside only, 2 – outside; default: 1, 1.
MajorTicksLength	d	2	length of major ticks (in relative units to the diagonal plot length). default: 0.03, 0.03.
TicksX	d	M	where to place ticks (x-axis)
TicksY	d	M	where to place ticks (y-axis)
MinorTicks	i	2	0 – automatic, 1 – no minor ticks, >1 – number of minor ticks intervals. default: 0.
MinorTicksStyle	i	2	0 – both directions, 1 – inside only, 2 – outside. default: 1, 1.
MinorTicksLength	d	2	length of major ticks (in relative units to the diagonal plot length). default: 0.015, 0.015.
GridLines	i	2	0 – no grid lines, 1 – dotted grid lines at major ticks, 2 – dashed grid lines at major ticks, 3 – dashed grid lines at major ticks, dotted at minor. default: 1, 1.
PlotRange	i	2	0 – automatic, 1 – user specified. default: 0, 0.
PlotRangeLimits[N]	d	2	lower and upper limits of plot range. default: 0, 1.
BinType	i	1	0 – no binning, 1 – averages values of bins 0,... Bin-1 into bin 0, 2 – averages values of bins 0,... Bin-1 into (Bin-1)/2, 3 – adds values of bins 0,... Bin-1 into bin 0, 4 – adds values of bins 0,... Bin-1 into (Bin-1)/2. default: 0.
Bin	i	1	number of bins. default: 1.

Parameter Name	Type	Dim	Description
MarkerType	i	1	0 – no marker, 1 – x-axis markers, 2 – y-axis markers, 3 – trace markers. default: 0.
Marker	d	2	Marker values. default: lower and upper plot range.
PlotStyle	i	16	0 – lines, 1 – symbols, 2 – lines and symbol. default: 0.
LineColor	i	16	–1 – automatic 0 – black, 1 – red, 2 – blue, 3 – green, 4 – yellow, 5 – brown, 6 – cyan, 7 – purple, default: –1 (0 on paper; 0,... 7 on screen)
LineStyle	i	16	–1 – automatic 0 – solid 1 – dashed 2 – dotted 3 – dash-dotted 4 – dash-dot-dotted 5 – dash-dash-dotted 6 – dash-dash-dot-dotted 7 – long dashed default: –1 (0 on screen, 0,... 7 on paper)
LineTickness	d	16	tickness of lines (in relative units to the diagonal plot length). default: 0.004.
SymbolShape	i	16	–1 – automatic 0 – circle, 1 – triangle up, 2 – square, 3 – triangle down, 4 – diamond, 5 – star, 6 – cross, 7 – slanted cross. default: –1 (0,... 7).
SymbolSize	d	1	size of symbols (in relative units to the diagonal plot length). default: 0.02.

Parameter Name	Type	Dim	Description
ErrorBars	i	16	0 – no error bars, 1 – symmetrical vertical error bars, 2 – symmetrical horizontal error bars, 3 – symmetrical vertical and horizontal error bars, 4 – asymmetrical vertical error bars, 5 – asymmetrical horizontal error bars, 6 – asymmetrical vertical and horizontal error bars. default: 0.
PlotNameA[M]	ch	1	Name of plot curve
PlotNameB[M]	ch	1	Name of plot curve, B channel (can be used to specify a second channel, e.g. B channel of transfer function, or B channel of cross-power spectrum).
PlotColumnA[M]	i	7	data columns: XY plot: X, Y, –, err1A, err2A, err1B, err2B; time series: –, Re Y, Im Y, err1A, err2A, err1B, err2B; histogram: X, Y, –, err1A, err2A, err1B, err2B; Bode plot: X, Re Y, Im Y, err1A, err2A, err1B, err2B. The error bar columns are used depending on the setting of ErrorBars. Default: 0, 1, 2, 3, 4, 5, 6.
PlotColumnB[M]	i	7	data columns; same as above, but for B channel.

Plotting packages which do not implement the complete set of above options have to implement as many as possible.

B.2.3 TIME SERIES

Time series containing raw channel information are stored by the channel name and two indices (measurement step and measurement point). Each of these data object contains the data of one channel only. The following table lists the parameters of a time series object:

Name	Type	Dim	Description
ObjectType	st	1	TimeSeries
Subtype	i	1	0 – normal time series in format (Y), 1 – down-converted time series in format (Y), 2 – averaged time series in format (Y), 3 – averaged time series in format (mean, std. dev., min., max., rms), 4 – normal time series in format (t,Y), 5 – down-converted time series in format (t,Y), 6 – averaged time series in format (t,Y), 7 – averaged time series in format (t, mean, std. dev., min., max., rms),
t0	ll	1	start time in GPS nsec.
dt	d	1	temporal spacing in sec.

Name	Type	Dim	Description
tp	d	1	pre-cursor time in sec. The time stamp of the first data point is: $t_0 - tp$.
f0	d	1	modulation frequency in Hz (type 2/5), trigger rate in Hz (type 1/4).
tf0	ll	1	start time for modulation signal in GPS nsec.
Decimation	i	1	overall decimation factor; default 1.
Decimation1	i	1	decimation factor before modulation (first stage)s; default 1.
DecmiationType	i	1	decimation filter type identifier; default 1 (see Table 1).
DecimationFilter	st	1	description of decimation filter.
DecimationDelay	d	1	delay introduced by the decimation filter in sec.
DelayTaps	i	1	number of taps in the time delay filter (in number of original samples).
TimeDelay	d	1	remaining time delay of time series in sec.
AverageType	i	1	0 – fixed number, 1 – running (exponential weight).
Averages	i	1	number of averages.
Channel	ch	1	channel name
N	i	1	number of points.
Unit	st	1	physical unit.
	f / zf	N 2×N 5×N 6×N	time series in format (Y), time series in format (t,Y), time series in format (mean, dev., min., max., rms), time series in format (t, mean, dev., min., max., rms).

Normally, raw data uses subtypes 0 or 1, whereas a time series of a result object uses subtype 2 or 3.

B.2.4 FFT AND (CROSS) POWER SPECTRUM

A data object which describes FFT and power spectra can contain multiple spectra. By convention the first one always describes the averaged spectrum. The following table lists the parameters associated with a power spectrum:

Name	Type	Dim	Description
Type	st	1	Spectrum
Subtype	i	1	0 – FFT in format (Y), 1 – power spectral density in format (Y), 2 – cross-power spectrum in format (Y), 3 – coherence in format (Y), 4 – FFT in format (f, Y), 5 – power spectral density in format (f,Y), 6 – cross-power spectrum in format (f,Y), 7 – coherence in format (f,Y).
f0	d	1	start frequency in Hz.
df	d	1	frequency spacing in Hz.
t0	ll	1	start time in GPS nsec.
dt	d	1	temporal spacing in sec (only useful for averaged power spectrum).
BW	d	1	resolution bandwidth
Window	i	1	0 – uniform (no window), 1 – Hanning, 2 – Flat-top, 3 – Welch, 4 – Bartlett, 5 – BMH, 6 – Hamming, 7 – Kaiser.
AverageType	i	1	0 – fixed number, 1 – running (exponential weight).
Averages	i	1	number of averages (only useful for power spectra).
ChannelA	ch	1	channel name
ChannelB[M]	ch	1	2nd channel name(s) for cross-power spectra
N	i	1	number of points.
M	i	1	number of spectra.
Unit	st	1	physical unit.
	f zf	M×N +N M×N +N	power spectral density/coherence in format (Y), power spectral density/coherence in format (f,Y), FFT/cross spectrum in format (Y), FFT/cross spectrum in format (f,Y).

Diagnostics tests will generally use subtype 1 for a power spectrum, subtype 2 for the cross-spectra and subtype 3 for the coherence. Power spectra are typically stored individually ($M = 1$), whereas cross-spectrum and coherence are stored as a set of spectra with one A channel and multiple B channels.

B.2.5 TRANSFER FUNCTION AND COHERENCE

A transfer function object can contain multiple transfer functions of the same two measurement points. By convention the first one is the average of the following ones. The following list presents the associated parameters of a transfer function:

Name	Type	Dim	Description
ObjectType	st	1	TransferFunction
Subtype	i	1	0 – transfer function B/A in format (Y), 1 – transfer function A in format (Y), 2 – coherence B/A in format (Y), 3 – transfer function B/A in format (f,Y), 4 – transfer function A in format (f,Y), 5 – coherence B/A in format (f, Y).
f0	d	1	start frequency in Hz.
df	d	1	frequency spacing in Hz.
t0	ll	1	start time in GPS nsec.
BW	d	1	measurement bandwidth in Hz.
Window	i	1	0 – uniform (no window), 1 – Hanning, 2 – Flat-top, 3 – Welch, 4 – Bartlett, 5 – BMH, 6 – Hamming, 7 – Kaiser. default: 1
AverageType	i	1	0 – fixed number, 1 – running (exponential weight).
Averages	i	1	number of averages.
ChannelA	ch	1	name of A channel
ChannelB[M]	ch	1	name(s) of B channel
N	i	1	number of points.
M	i	1	number of transfer/coherence functions.
	zf f	M×N +N M×N +N	transfer function in format (Y), transfer function in format (f,Y), coherence in format (Y), coherence in format (f,Y).

Diagnostics tests will generally use subtype 3 for a transfer function and subtype 5 for the coherence. Typically, they are both stored as a set of functions with one A channel and multiple B channels.

B.2.6 LIST OF COEFFICIENTS

A sine response measurement can yield multiple transfer coefficients which are stored in two dimensional arrays. One of the dimension always represents the multiple measurement points, whereas the other dimension may represent the multiple excitation points, the harmonic order or the modulation product terms.

Name	Type	Dim	Description
ObjectType	st	1	Coefficients
Subtype	i	1	0 – transfer coefficients in format (Y), 1 – harmonic coefficients in format (Y), 2 – intermodulation product in format (Y), 3 – coherence coefficients in format (Y), 4 – transfer coefficients in format (f, Y), 5 – harmonic coefficients in format (f, Y), 6 – intermodulation product in format (f, Y), 7 – coherence coefficients in format (f, Y), 8 – transfer matrix in format (Y).
t0	ll	1	start time in GPS nsec.
BW	d	1	measurement bandwidth in Hz.
AverageType	i	1	0 – fixed number, 1 – running (exponential weight), 2 – running (accumulative).
Averages	i	1	number of averages.
ChannelA[M']	ch	1	channel name corresponding to M'-th readback point.
ChannelB[N]	ch	1	channel name corresponding to N-th detection point.
N	ll	1	number of detection points (A + B channels, except subtype 8 which only counts B channels)

Name	Type	Dim	Description
M	ll	1	number of frequency points: f_1, f_2, f_M: transfer/coherence coefficients, 0, f, 2 f, 3 f, M f: harmonic coefficients (the 0 frequency is used for storing the total harmonic distortion), f_1, f_2, f_1 - f_2 , f_1 + f_2: intermodulation product
Unit[N]	st	1	physical unit.
	zf zf zf f zf zf zf zf f	M×N M× (N+1)	transfer coefficients in format (Y), harmonic coefficients in format (Y), intermodulation product in format (Y), coherence coefficients in format (Y), transfer matrix in format (Y), transfer coefficients in format (f, Y), harmonic coefficients in format (f, Y), intermodulation product in format (f, Y), coherence coefficients in format (f, Y).

B.2.7 MEASUREMENT VALUES

Measurement values which do not fit into one of the above category can be stored in a table:

Name	Type	Dim	Description
ObjectType	st	1	MeasurementTable
t0	ll	1	start time in GPS nsec.
TableLength	i	1	length of table
Name[N]	st	1	name of measurement variable
Unit[M]	st	1	physical unit.
Description[M]	st	1	description of measurement
ValueType[M]	st	1	type of value: number or complex
	zf	M	measurement values

B.3 DIAGNOSTICS TESTS

B.3.1 SINE RESPONSE, HARMONIC DISTORTION AND TWO-TONE INTERMODULATION TESTS

Name	Type	Dim	Description
ObjectType	st	1	TestParameter
Subtype	st	1	SineResponse
MeasurementTime	d	2	measurement time at each frequency; first value in sec, second value in cycles. The smaller one is taken and rounded up to the next cycle. Negative values are ignored. default: 100, 10.
AverageType	i	1	0 – fixed number, 1 – running (exponential weight), 2 – running (accumulative). default: 0.
Averages	i	1	number of averages; default is 1.
SettlingTime	d	1	settling time; specified in fraction of the total measurement time; default is 0.1.
StimulusChannel[M]	ch	1	name of stimulus channel.
StimulusReadback[M]	ch	1	if defined, this is the readback of the stimulus channel, otherwise, it is assumed the readback channel is identical to the stimulus channel. A stimulus readback can be explicitly disabled by specifying a “!” as the stimulus readback channel.
StimulusFrequency[M]	d	1	frequency of stimulus channel in Hz. default: 100.
StimulusAmplitude[M]	d	1	amplitude of stimulus channel. default: 0.
StimulusOffset[M]	d	1	offset of stimulus channel. default: 0.
StimulusPhase[M]	d	1	phase of stimulus signal in rad; the phase is relative to the last 0:00UTC. default: 0.
MeasurementChannel[N]	ch	1	measurement channel.

Name	Type	Dim	Description
HarmonicOrder	i	1	indicates the highest harmonic order of interest; when set to a value larger than 1 (default), it prevents that the data to be decimate to a rate below $2 * \max(\text{StimulusFrequency}) * \text{HarmonicOrder}$. A value of zero or smaller will prevent any decimation.
Window	i	1	0 – uniform (no window), 1 – Hanning, 2 – Flat-top, 3 – Welch, 4 – Bartlett, 5 – BMH, 6 – Hamming, 7 – Kaiser. default: 1
FFTResult	b	1	if true (default), calculates averaged 1024 point FFTs of each measurement channel as part of the result.

If there is only one stimulus channel, the result will automatically include a harmonic analysis of the first stimulus channel at every measurement point. Similarly, if there are exactly two stimulus channels of different frequencies, the result will include a two-tone intermodulation analysis at every measurement point. In all cases the result will include a sine response analysis of every stimulus frequency at every measurement point.

B.3.2 SWEPT SINE TESTS

Name	Type	Dim	Description
ObjectType	st	1	TestParameter
Subtype	st	1	SweptSine
SweepType	i	1	0 – linear, 1 – log, 2 – user supplied frequency points, 3 – user supplied frequency/amplitude points default: 1.
SweepDirection	i	1	0 – upwards, 1 – downwards (default).
StartFrequency	d	1	start frequency of swept sine in Hz. default: 1.
StopFrequency	d	1	stop frequency of swept sine in Hz. default: 1000.
NumberOfPoints	i	1	number of frequency steps; default is 61.

Name	Type	Dim	Description
SweepPoints	d	M	user supplied frequency and amplitude points: if SweepType is equal 2 the points consist of only frequency values; if SweepType is 3, they consist of frequency-amplitude pairs.
AChannels	i	1	Number of A channels; default 0. If this value is non-zero, the specified amount of channels starting with the first measurement channel are also considered A channels. By default only the excitation channel is considered an A channel. For each A channels, the swept sine test will calculate both the transfer function and the coherence between them and any other (B) channel. A negative number will automatically include all channels.
Averages	i	1	number of averages; default is 1.
MeasurementTime	d	2	measurement time at each frequency; first value in sec, second value in cycles. The smaller one is taken and rounded up to the next cycle. Negative values are ignored. default: 100, 10.
SettlingTime	d	1	settling time at each frequency step; specified in fraction of the total measurement time; default is 0.1.
StimulusChannel	ch	1	name of stimulus channel.
StimulusReadback	ch	1	if defined, this is the readback of the stimulus channel, otherwise, it is assumed the readback channel is identical to the stimulus channel.
StimulusAmplitude	d	1	amplitude of stimulus channel. default: 0.
MeasurementChannel[N]	ch	1	measurement channels; must be at least two. By default the first channel is the A channel.
HarmonicOrder	i	1	indicates the highest harmonic order of interest; when set to a value larger than 1 (default), it prevents that the data to be decimate to a rate below $2 * \max(\text{StimulusFrequency}) * \text{HarmonicOrder}$. A value of zero or smaller will prevent any decimation.

Name	Type	Dim	Description
Window	i	1	0 – uniform (no window), 1 – Hanning, 2 – Flat-top, 3 – Welch, 4 – Bartlett, 5 – BMH, 6 – Hamming, 7 – Kaiser. default: 1
FFTResult	b	1	if true, calculates averaged 1024 point FFTs of each measurement channel for each frequency step as part of the result; default is false.

B.3.3 FOURIER TESTS

Name	Type	Dim	Description
ObjectType	st	1	TestParameter
Subtype	st	1	FFT
StartFrequency	d	1	start frequency of FFT in Hz. default: 0.
StopFrequency	d	1	stop frequency of FFT in Hz; internally always rounded up so that the frequency span is a power of 2. default: 1000.
BW	d	1	bandwidth in Hz; always rounded to the closest power of 2. default: 1.
Overlap	d	1	overlap of FFT windows (only useful when averaging); $t(i) = t_0 + i * (1 - \text{Overlap}) * dt$. default: 0.5.
Window	i	1	0 – uniform (no window), 1 – Hanning, 2 – Flat-top, 3 – Welch, 4 – Bartlett, 5 – BMH, 6 – Hamming, 7 – Kaiser. default: 1

Name	Type	Dim	Description
AChannels	i	1	Number of A channels; default 0. If this value is non-zero, the specified amount of channels starting with the first measurement channel are considered A channels. For A channels, the fft test will calculate both the cross-correlation and the coherence between them and any other channel. Any stimulus readback channel is automatically considered an A channel. A negative number will automatically include all channels.
AverageType	i	1	0 – fixed number, 1 – running (exponential weight), 2 – running (accumulative). default: 0.
Averages	i	1	number of averages; default is 10.
SettlingTime	d	1	settling time; specified in fraction of the measurement time; default is 0.
StimulusType[M]	i	1	0 – no signal, 1 – sine wave, 2 – square wave, 3 – ramp, 4 – triangle, 5 – impulse, 6 – constant amplitude, 7 – normally distributed noise, 8 – uniformly distributed noise, 9 – arbitrary waveform, 10 – linear amplitude/frequency sweep, 11 – logarithmic amplitude/frequency sweep. default: 0.
StimulusChannel[M]	ch	1	name of trigger channel.
StimulusReadback[M]	ch	1	if defined, this is the readback of the stimulus channel, otherwise, it is assumed the readback channel is identical to the stimulus channel. A stimulus readback can be explicitly disabled by specifying a “!” as the stimulus readback channel.
StimulusFrequency[M]	d	1	frequency of trigger signal; default is 0, i.e. $1/\text{MeasurementTime}$. describes the center frequency for sweeps and band-limited noise signals; describes the sampling frequency for arbitrary waveforms.
StimulusAmplitude[M]	d	1	amplitude of trigger signal; default 0. describes the center amplitude for sweeps.
StimulusOffset[M]	d	1	offset of trigger signal; default 0. (describes the delay in sec for impulse).

Name	Type	Dim	Description
StimulusPhase[M]	d	1	phase of stimulus signal in rad; the phase is relative to the last 0:00UTC. default: 0. (describes the duration in sec for impulse and sweeps).
StimulusRatio[M]	d	1	ratio between high and the full period of a square wave signal; default is 0.5.
StimulusFrequencyRange[M]	d	1	describes the frequency range for sweeps and band-limited noise signals; default 10000.
StimulusAmplitudeRange[M]	d	1	describes the amplitude range for sweeps; default 0.
StimulusPoints[M]	f	K	array describing an arbitrary waveform.
MeasurementChannel[N]	ch	1	measurement channel.

B.3.4 TIME SERIES MEASUREMENTS AND TRIGGER RESPONSE TESTS

Name	Type	Dim	Description
ObjectType	st	1	TestParameter
Subtype	st	1	TimeSeries
MeasurementTime	d	1	measurement time in sec. To allow simple averaging this value is rounded to the next multiple of the sample period of the measurement channel with the lowest sampling rate. The sampling rate of a channel can be reduced by specifying a bandwidth.
PreTriggerTime	d	1	measurement time before trigger is applied (given as a ratio); default is 20% of the total measurement time.
DeadTime	d	1	dead time; specified in fraction of the measurement time; default is 0.
SettlingTime	d	1	settling time; specified in fraction of the measurement time; default is 0.
BW	d	1	required signal bandwidth in Hz; always rounded up to the closest power of 2. This value is used to determine the decimation factor. default: 10000.
IncludeStatistics	b	1	If true (default), the resulting time series includes the mean, the standard deviation, the minimum, the maximum and the root-mean-square value. If false, only the mean is included.
Averages	i	1	number of triggers/averages. 1 – single trigger response (no average), >1 – periodic trigger response (with average), default: 10.

Name	Type	Dim	Description
AverageType	i	1	0 – fixed number, 1 – running (exponential weight), 2 – running (accumulative). default: 0.
StimulusType[M]	i	1	0 – no signal, 1 – sine wave, 2 – square wave, 3 – ramp, 4 – triangle, 5 – impulse, 6 – constant amplitude, 7 – normally distributed noise, 8 – uniformly distributed noise, 9 – arbitrary waveform, 10 – linear amplitude/frequency sweep, 11 – logarithmic amplitude/frequency sweep. default: 0.
StimulusChannel[M]	ch	1	name of trigger channel.
StimulusReadback[M]	ch	1	if defined, this is the readback of the stimulus channel, otherwise, it is assumed the readback channel is identical to the stimulus channel. A stimulus readback can be explicitly disabled by specifying a “!” as the stimulus readback channel.
StimulusFrequency[M]	d	1	frequency of trigger signal; default is 0, which is interpreted as 1/(MeasurementTime). describes the center frequency for sweeps and band-limited noise signals; describes the sampling frequency for arbitrary waveforms.
StimulusAmplitude[M]	d	1	amplitude of trigger signal; default 0. describes the center amplitude for sweeps.
StimulusOffset[M]	d	1	offset of trigger signal; default 0. (describes the delay in sec for impulse).
StimulusPhase[M]	d	1	phase of stimulus signal in rad; the phase is relative to the trigger point. default: 0. (describes the duration in sec for impulse and sweeps).
StimulusRatio[M]	d	1	ratio between high and the full period of a square wave signal; default is 0.5.
StimulusFrequencyRange[M]	d	1	describes the frequency range for sweeps and band-limited noise signals; default 10000.
StimulusAmplitudeRange[M]	d	1	describes the amplitude range for sweeps; default 0.
StimulusPoints[M]	f	K	array describing an arbitrary waveform.
MeasurementChannel[N]	ch	1	measurement channel.

B.3.5 RANDOM STIMULUS RESPONSE TESTS

TBD.

B.4 XML CONVENTIONS

The file format for saving a diagnostics test is the LIGO lightweight file format which is based on XML. Data fields of parameters are ASCII encoded; in case the parameter uses a list of values they are comma delimited. Data fields of data objects are save as binary arrays following the C convention for storing rows and columns and using a big-endian representation. These binary data fields are appended to the XML file, so that a diagnostics test can be stored in a single file.

APPENDIX C SOFTWARE MODULES

A fair amount of documentation is written directly into the header files of the C and C++ modules. This documentation can be converted into html web pages using the doc++ program. The web pages can be reached at www.ligo-wa.caltech.edu/gds. This appendix is intended to give an overview of the existing software modules and to explain their main functions. A detailed description of the constants, types, macros, objects, routines and their parameters is found in the web pages.

C.1 OVERALL STRUCTURE

The source tree of the diagnostics test software is divided into sections (sub directories) representing groups of modules. A typical section contains all the modules associated with a software interface or category. The following sections are implemented:

Section	Description	Modules
src/algorithm	analysis algorithms: FFT, swept sine, etc.	decimate – decimation gdsrand – random number generation gdssigproc – general purpose signal processing sineanalyze – sine amplitude determination
src/awg	arbitrary waveform generator, excitation engine	awg – excitation generator and manager awgapi – API to the excitation engine awgfunc – general purposes awg functions awgtype – types used by the excitation engine awg_server – rpc server of excitation engine excitation – diagnostics test interface to the awg rawgapi – rpc (remote procedure call) interface
src/cmd	command line interface	cmdline – command line interpreter gdscmd – cmd line interface to diagnostics kernel gdsmsg – message interface for sending commands gdsmsg_server – rpc server for command messages rgdsmsg – rpc interface
src/daq	interface to the data acquisition system	gdschannel – interface to channel information database gdsrtdd – interface to the network data server
src/diag	diagnostics tests	diagclass – basic class for diagnostics tests diagnames – names used by tests repeat – repeats a test sineresponse – sine response test sweptsine – swept sine test testiter – basic object for test iterations testorg – manages all diagnostics tests

Section	Description	Modules
src/drv	hardware drivers	cobox – driver for an ethernet-to-RS232 converter ds340 – driver for SR DS340 signal generators gdsdac – driver for ICS115 digital-to-analog converter gpsclk – driver for the VME gps clocks hardware – parameters of the VME modules rmap – driver for reflective memory boards target – VME host parameters
src/prog	programs	chndump – dumps the channel database to screen chnsave – saves channel data to disk diag – main program gdsd – diagnostics kernel (daemon)
src/rmem	reflective memory interface	map – memory map of the reflective memory rmorg – macros for managing DCUs testpoint – API for selecting test points testpoint_server – test point manager rtestpoint – rpc interface
src/sched	scheduler	gdssched – scheduler gdssched_client – remote scheduler client interface gdssched_server – remote scheduler server interface gdsrsched – rpc interface
src/storage	storage objects for diagnostics tests	gdsdatum – hierarchical storage object diagdatum – diagnostics storage object rtddinput – storage interface to the network data server
src/test	test programs	miscellaneous
src/util	utilities	gdserr – common error log gdserrmsg – error messages gdsheartbeat – heartbeat interface and synchronization gdsmain – compiler directives for different hosts gdsmutex – mutual exclusion semaphore objects gdsprm – parameter file interface gdsstring – additional string functions gdstask – task/thread creation rpcinc – common rpc routines tconv – time conversion routines gdsutil – includes most utility modules

C.2 UTILITIES

The following sections list the main routines and objects of the more important software modules.

C.2.1 GDSERR

Function	Description
<code>gdsConsoleMessage</code>	prints a message to the gds console
<code>gdsErrorMessage</code>	prints an error message
<code>gdsError</code>	prints one of the predefined error conditions
<code>gdsWarningMessage</code>	prints a warning message
<code>gdsDebugMessage</code>	prints a debug message
<code>gdsDebug</code>	prints a debug message if the DEBUG is defined

C.2.2 GDSHEARTBEAT

Function	Description
<code>installHeartbeat</code>	installs a heartbeat interrupt (16 Hz clock derived from GPS)
<code>syncWithHeartbeat</code> <code>syncWithHeartbeatEx</code>	synchronizes the program execution to the next heartbeat

C.2.3 GDSMUTEX

Object	Description
<code>mutex</code>	mutex object
<code>recursivemutex</code>	mutex which can be called multiple times from the same thread
<code>readwritelock</code>	read/write lock semaphore
<code>semlock</code>	locks a semaphore/mutex as long during the scope of the object

C.2.4 GDSPRM

Function	Description
<code>findParamFileSection</code>	finds a section within a parameter file
<code>nextParamFileSection</code>	finds the next section
<code>getParamFileSection</code>	gets a section
<code>findParamSectionEntry</code>	finds a parameter section entry
<code>nextParamSectionEntry</code>	finds the next entry
<code>getParamSectionEntry</code>	gets the entry

Function	Description
loadParamSectionEntry	loads a parameter
loadBoolParam	loads a boolean parameter
loadIntParam	loads an integer parameter (int)
loadNumParam	loads a numerical value (unsigned long)
loadFloatParam	loads a floating point number (double)
loadStringParam	loads a string parameter

C.2.5 GDSSTRING

Function	Description
gds_strcasecmp	compares two strings ignoring the case
strend	goes to the end of a string
strecpy	copies a string and returns the end of the resulting string
chnIsValid	returns true if a correctly formatted channel name is supplied
chn...	miscellaneous channel name handling functions

C.2.6 GDSTASK

Function	Description
taskCreate	creates a new task/thread

C.2.7 RPCINC

Function	Description
rpcGetHostaddress	returns the host address
rpcGetLocalAddress	gets the address of the local machine
rpcGetClientAddress	gets the rpc client address
rpcInitializeServer	initializes an rpc server
rpcRegisterService	registers an rpc service
rpcStartServer	starts an rpc server
rpcStartCallbackService	starts an rpc callback service
rpcStopCallbackService	terminates the rpc callback service
rpcResgisterCallback	registers a callback service
rpcProbe	test if an rpc service exists

C.2.8 TCONV

Function	Description
TAInow	returns the time in GPS seconds
TAltoUTC UTCtoTAI	converts between universal coordinate time and gps seconds

C.3 ALGORITHMS

C.3.1 DECIMATE

Function	Description
decimate	filter decimation stage
zoom	down-conversion

C.3.2 GDSRAND

Function	Description
urand_r	MT safe uniformly distributed random number generator
urandv_r	uniformly distributed random vector generator
nrand_r	normally distributed random number generator
nrandv_r	normally distributed random vector generator

C.3.3 GDSSIGPROC

Function	Description
DotProd	vector dot product
Mean	vector mean
Mixdown	complex down-conversion

C.3.4 SINEANALYZE

Function	Description
sineAnalyze	determines the amplitude and phase of a sine wave
sweptSineNpts	calculates the number of necessary data points

C.4 ARBITRARY WAVEFORM GENERATOR

C.4.1 AWG

Function	Description
initAwg	initializes the arbitrary waveform generator
getIndexAWG	gets an unused slot in the awg
releaseIndexAWG	frees an awg slot
resetAWG	resets an awg slot
configAWG	configures an awg from a parameter file
showAWG	displays the current state information of an awg slot
processAWG	calculates the waveform
disableAWG	disables a slot
enableAWG	enables a slot
addWaveformAWG	adds a waveforms to a slot
setWaveformAWG	sets an arbitrary waveform vector
queryWaveformAWG	returns the current waveforms
checkConfigAWG	checks the configuration
resetAllAWG, configAllAWG, showAllAWG, processAllAWG, disableAllAWG, enableAllAWG	same as corresponding functions above but for all awg slots
getStatisticsAWG	get the real-time performance statistics of an awg

C.4.2 AWGAPI

Function	Description
awg_client	installs the awg client interface
awgSetChannel	reserves a slot of an awg
awgRemoveChannel	frees the slot
awgAddWaveform	adds waveforms (signals)
awgSetWaveform	sets an arbitrary waveform vector
awgQueryWaveform	returns the current waveforms
awgReset	resets an awg

Function	Description
awgStatistics	returns the real-time performance statistics of an awg
awgShow	shows the status information of an awg
awgCommand	command interpreter for arbitrary waveform generators

C.4.3 AWGFUNC

Function	Description
normPhase	normalizes the phase to be between 0 and 2π
productLog	calculates the inverse of $z(w)=w \exp(w)$
awgSignal	calculates a waveform function
awgPhaseIn	phase-in function
awgPhaseOut	phase-out function
awgSweepOut	phase-out for frequency sweep
awgSweepComponents	defines a frequency sweep waveform component
awgPeriodicComponent	defines a waveform component of a periodic waveform
awgIsValidComponent	tests if waveform component is valid
awgSortComponents	sorts awg components according to their start time

C.4.4 AWG_SERVER

Function	Description
awg_server	starts the rpc services for an arbitrary waveform generator

C.4.5 EXCITATION

TBD.

C.5 COMMAND LINE INTERFACE

C.5.1 CMDLINE

Object / Method	Description
commandline	command line object
operator !	returns true if finished
operator ()	calls the command line interpreter

C.5.2 GDSCMD

Function	Description
gdsCmdInit	Initializes the diagnostics kernel
gdsCmdFini	terminates the diagnostics kernel
gdsCmd	executes a diagnostics command
gdsCmdNotifyHandler	installs a command notification handler
cmdNotification	sends a notification back to the command line interface

C.5.3 GDSMSG

Function	Description
gdsMsgOpen	opens a communication channel to the diagnostics kernel
gdsMsgClose	closes the communication channel
gdsMsgSend	sends a message to the diagnostics kernel
gdsMsgInstallHandler	installs a message callback handler for receiving notifications

C.5.4 GDSMSG_SERVER

Function	Description
gdsmsg_server	starts the rpc message server of the diagnostics kernel

C.5.5 GDSMSG_SOCKETS

Function	Description
gdsmsg_sockets	starts the TCP/IP message server of the diagnostics kernel

C.6 DATA ACQUISITION INTERFACE

C.6.1 GDSCHANNEL

Function	Description
gdsChannelInfo	obtains channel information
gdsChannelListLen	returns the number of channels in the database
gdsChannelList	returns a list of all channels

C.6.2 GDSRTDD

Function	Description
gdsSubscribeData	subscribes a channel to the network data server
gdsUnsubscribeData	unsubscribes the channel
gdsGetNewdata	gets new channel data
gdsGetData	gets channel data for a specific time interval

C.7 DIAGNOSTICS TEST UTILITIES**C.7.1 DIAGNAMES****C.7.2 DIAGORG****C.7.3 TESTENV****C.7.4 TESTSYNC****C.8 DIAGNOSTICS TEST SUPERVISORY****C.8.1 SUPERVISORY****C.8.2 STDSUPER**

TBD.

C.9 DIAGNOSTICS TEST ITERATORS**C.9.1 TESTITER****C.9.2 REPEAT****C.9.3 SCANITER****C.9.4 FINDITER****C.10DIAGNOSTICS TESTS****C.10.1DIAGTEST**

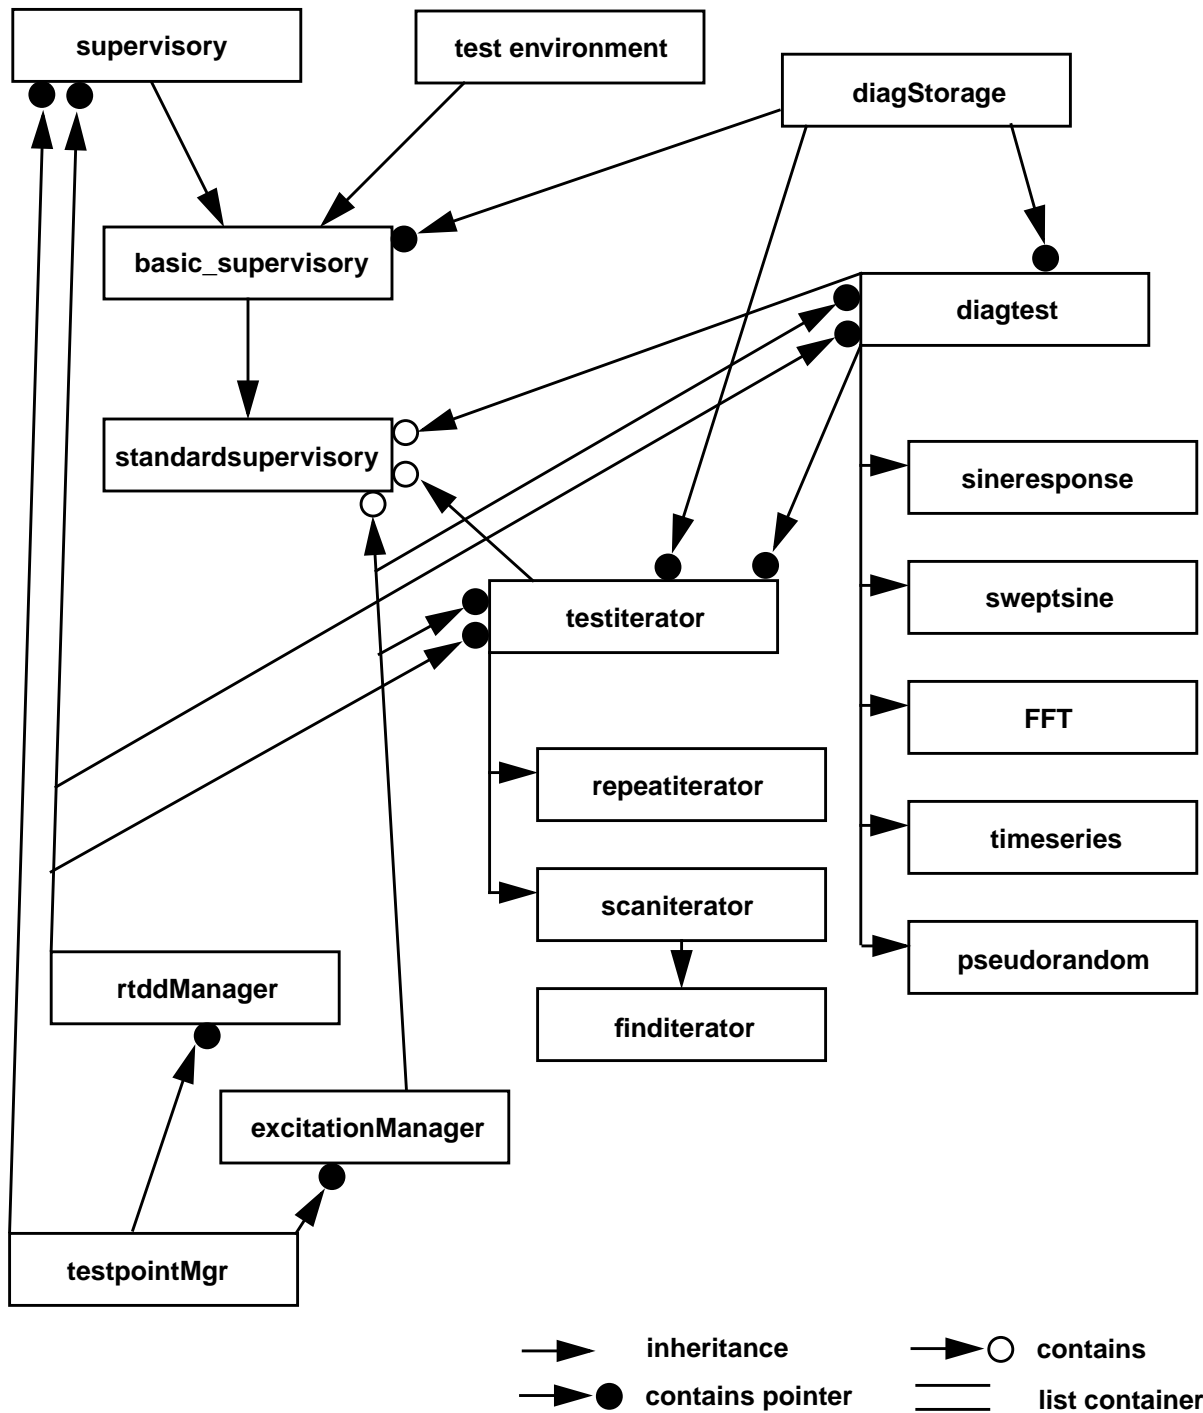


Figure 13: Class hierarchy of the diagnostics supervisory, the diagnostics tests and the test iterators.

C.10.2SINERESPONSE**C.10.3SWEPTSINE****C.10.4FOURIER****C.10.5TIMESERIES****C.10.6PSEUDORAN****C.11HARDWARE DRIVERS**

The layout of the excitation engine VME modules can be found in the hardware and target header files.

C.11.1COBOX

Function	Description
openCobox	opens a socket connection to a cobox (ethernet-to-RS232 converter)

C.11.2DS340

Function	Description
connectSerialDS340	connects a DS340 through a serial port
connectCoboxDS340	connects a DS340 through a cobox
resetDS340	resets a DS340
isDS340Alive	tests if the DS340 is alive
pingDS340	tests if the DS340 is connected and powered up
setDS340 uploadDS340Block	sets the configuration of a DS340
getDS340 downDS340Block	gets the configuration
sendWaveDS340	uploads an arbitrary waveform vector to a DS340
sendResetDS340	sends a reset signal
sendClearDS340	sends a clear signal
sendTriggerDS340	sends a trigger signal

C.11.3GDSDAC

Function	Description
dacInit	initializes the ICS115 board
dacRestart	restarts the ICS115 and re-synchronizes it with the GPS clock
dacCopyData	copies data to the ICS115

Not yet implemented.

C.11.4GPSCLK

Function	Description
gpsBaseAddress	returns the VME base address of the GPS clock
gpsInit	initializes the board
gpsSynInfo	returns the synchronization status
gpsTimeNow	returns the current time in GPS nsec
gpsTime	converts the native time format of the board into GPS nsec
gpsNativeTime	returns the current time in the native format
gpsMicroSec	returns the micro seconds only
gpsInfo	returns GPS information
gpsHeartbeatInstall	Installs a 16Hz GPS synchronized interrupt
gpsHeartbeatHealth	monitors the health of the heartbeat interrupt

C.11.5RMAPI

Function	Description
rmInit	initializes the reflective memory board
rmBaseAddress	returns the base address of the reflective memory region
rmBoardAddress	returns the base address of the reflective memory board
rmBoardSize	returns the memory size supported by the board
rmLED	turns the LED on and off
rmResetNode	resets a reflective memory node
rmInt	sends an interrupt to a reflective memory node
rmCheck	check if a memory region is accessible
rmRead	reads from reflective memory
rmWrite	writes to reflective memory

C.12 REFLECTIVE MEMORY ORGANIZATION

The memory map of the reflective memory and set of macros to handle parameters of data collection units can be found in the header files of the map and rmorg modules.

C.12.1 TESTPOINT

Function	Description
testpoint_client	initializes the test point client interface
tpRequest	selects a set of test points
tpClear	clears test points
tpQuery	returns the active test points
tpGetIndexDirect	returns the test point index on VME systems
tpIsValid	tests if a test point is valid
tpAddr	returns the address of a test point in reflective memory
tpCommand	command line interface to the test point manager

C.12.2 TESTPOINT_SERVER

Function	Description
testpoint_server	starts the rpc services of the test point manager

C.13 SCHEDULER

C.13.1 GDSSCHED

Function	Description
createScheduler	creates a scheduler
closeScheduler	closes a scheduler
scheduleTask	adds a new task to the scheduler
getScheduledTask	obtains information about a scheduled task
removeScheduledTask	removes a task from a scheduler
waitForSchedulerToFinish	waits until all tasks have been scheduled and finished
setScheduleTag	sets a synchronization tag

C.13.2GDSSCHED_CLIENT

Function	Description
createRemoteScheduler	creates a new local scheduler for remote use
createBoundScheduler	creates a new scheduler on a remote machine

C.13.3GDSSCHED_SERVER

Function	Description
registerSchedulerClass	registers a task which can be scheduler from remote
runSchedulerService	start the rpc services of a remote scheduler manager

C.14 STORAGE OBJECTS

The class hierarchy of the diagnostics storage interface is shown in Fig. 14. There are two main class categories: (i) classes which contain data (inherited from `gdsDatum`) and (ii) classes which manage the access to the data (inherited from `diagObjectName`). The main diagnostics storage object is of type `diagStorage`, it is inherited from a generic storage object, `gdsStorage`, and contains all access classes. This main object is able to set and get variables by name from the generic storage object using the access classes to make sure the variable names are valid test data object and valid test parameters. The storage object is organized implementing two hierarchical levels: the main storage objects contains both data objects and global parameter objects; whereas data objects can contain their own private parameter objects. Thus, the generic storage object, `gdsStorage` inherits from the data object, `gdsDataObject`, and also contains a list of data objects.

Access classes for data objects are inherited from `diagObject` which itself is inherited from `diagObjectName`. The `diagObject` contains a list of access classes for the parameters which are associated with this data object. Data objects which contain results and data objects which describe a diagnostics test are inherited from `diagMultipleObject`. The `diagMultipleObject` manages a list of access classes which are used depending on the type of the data object.

C.14.1GDSDATUM

Object / Method	Description
<code>gdsDatum</code>	object which stores data, can be multi-dimensional.
<code>gdsNamedStorage</code>	object with a name.
<code>gdsNamedDatum</code>	object which stores data and has a name.
<code>gdsDataReference</code>	object which manages data references to memory mapped files.
<code>storage_ptr</code>	auto pointer object which manages pointers to parameter and data objects.

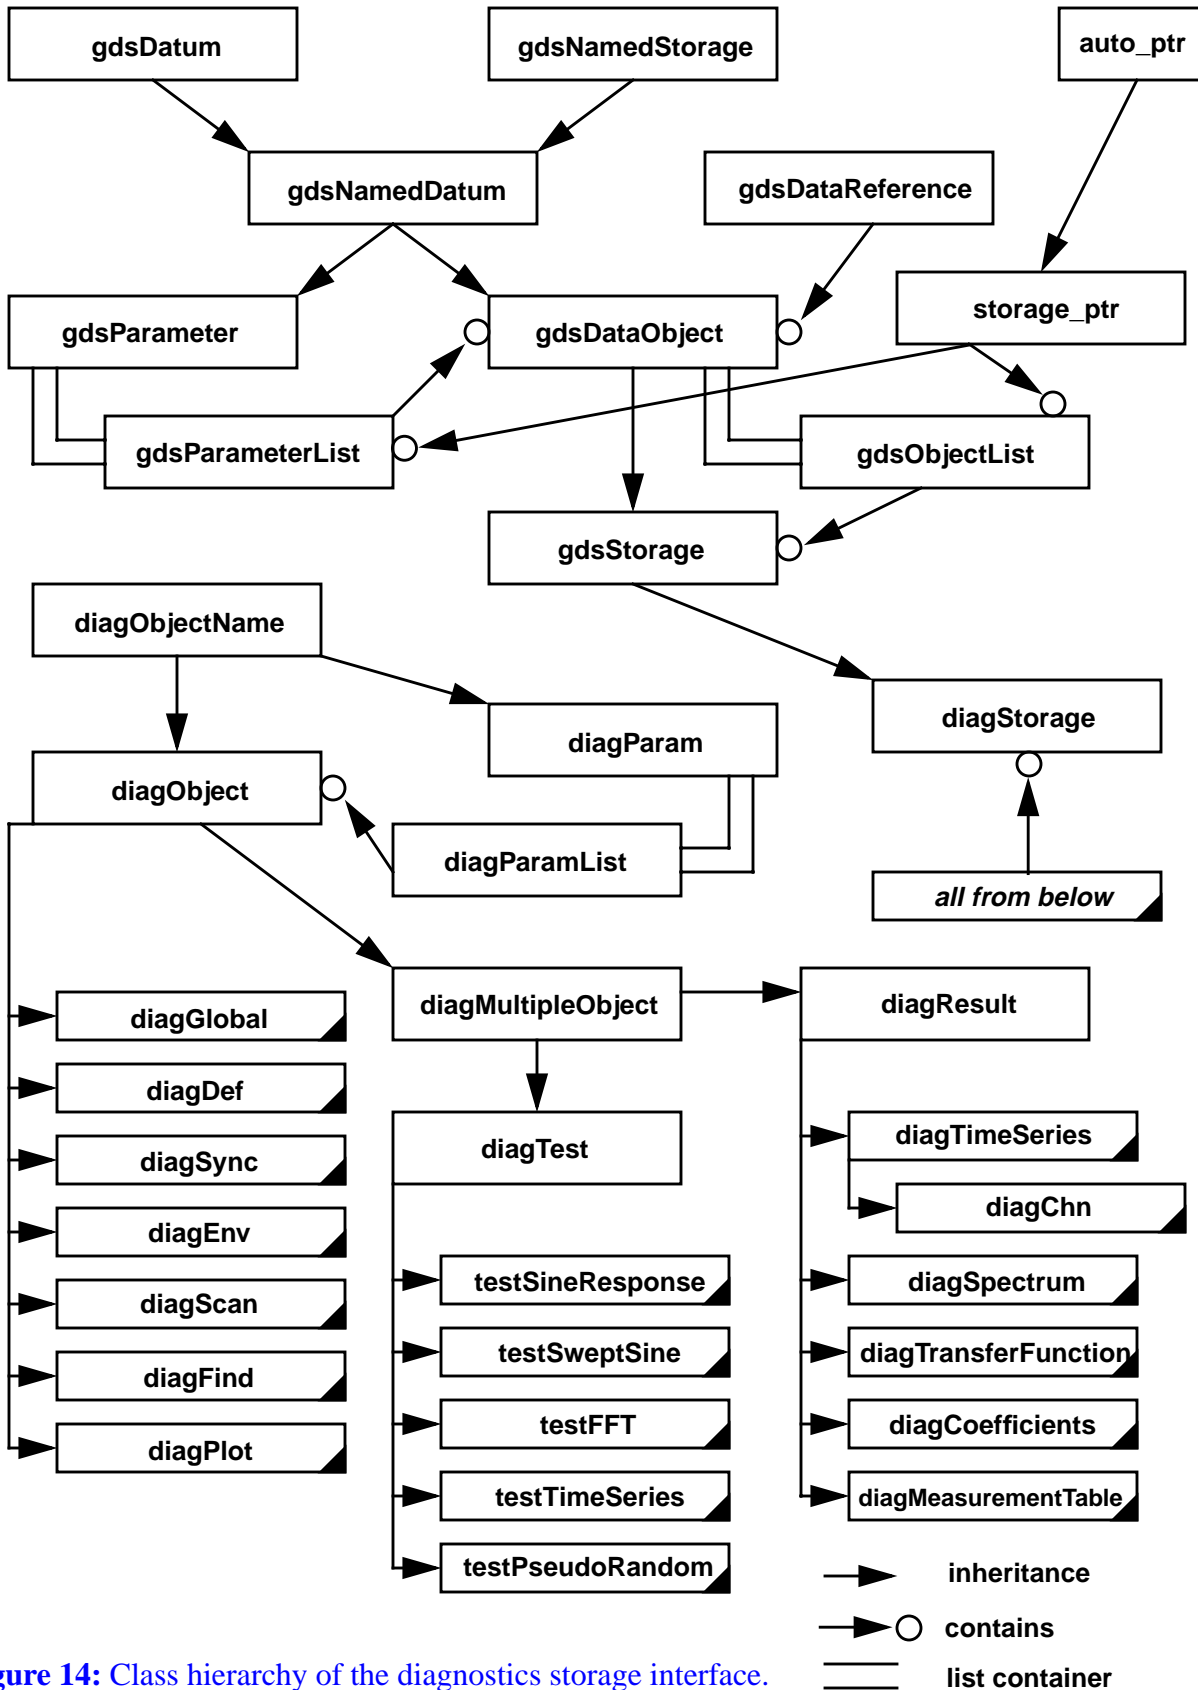


Figure 14: Class hierarchy of the diagnostics storage interface.

Object / Method	Description
gdsParameter	parameter object.
gdsDataObject	data object; contains a list of parameter objects.
gdsStorage	generic storage object; contains a list of data objects.

C.14.2DIAGDATUM

Object / Method	Description
diagObjectName	generic access class, handles names, dimensions, indices, data types and access writes.
diagParam	generic access class for a parameter object.
diagObject	generic access class for data objects.
diagGlobal, diagDef, diagSync, diagEnv, diagScan, diagFind, diagPlot	access class for global parameters, default parameters, synchronization variables, environment settings, scan parameters, optimization parameters and plot settings, respectively.
diagMultipleObject	generic access class for a data object with multiple possible configurations.
diagTest	access class for the diagnostic test object.
testSineResponse, testSweptSine, testFFT, testTimeSeries, testPseudoRandom	access class for sine response test, swept sine test, FFT test, time series measurement, and pseudo random response test, respectively.
diagResult	access class for a result object.
diagTimeSeries, diagChn diagSpectrum, diagCoefficients, diagMeasurementTable	access class for time series vectors, channel data, FFT spectra, measurement coefficients, and measurement tables, respectively.
diagStorage	diagnostics storage object; manages parameters and data objects which can have their own parameters. Uses the access classes to make sure that only valid parameter and data objects are stored.

C.14.3RTDDINPUT

Object / Method	Description
rtddcallback	class implementing a callback method for the network data server API
gdsRTDDchannel	class for reading channel data from the network data server, it manages a single channel, knows how to handle data partitions and knows where to store the read data in the diagnostics storage object.
partition	class which describes a channel data partition
gdsRTDDinput	class which manages a list of gdsRTDDchannel objects, i.e. class which handles the input from the network data server.

C.15 PROGRAMS

C.15.1 CHNDUMP

This is a utility program to dump all channel information records to the standard output.

C.15.2 DIAG

This program is the interface to the diagnostic system, it can invoke either the command line interface or the graphical user interface. For a more detailed discussion see Sections 2.2 – 2.4.

C.15.3 GDSD

This is the diagnostics kernel. See Appendix A.2 on how to set it up.

C.15.4 LIBGDS.SO

This is the dynamic link library which implements most of the diagnostics kernel. It is used by the ‘diag’ program if a connection to a local kernel is established.