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

Technical Note LIGO-T960031 - 00-E 3/7/96

CDS Online Diagnostic and Readout Functions

M.E. Zucker

Distribution of this draft:

Data Analysis working group

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (818) 395-2129
Fax (818) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

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

1 INTRODUCTION

This note describes desirable realtime (or almost realtime) online operation and diagnostic functions which may be implemented by LIGO CDS Data Acquisition, Remote Diagnostics and Operator Interface hardware and software. In order to concentrate on functionality without preconceived configuration bias, no explicit distinction is drawn between Remote Diagnostics and Data Acquisition subsystem roles. Any online *physics* data analysis is expressly excluded from the scope of this document; here we are strictly concerned with commissioning, operating and characterizing LIGO interferometers.

Functional descriptions are loosely based on the typical operation and display capabilities found (or sorely missed) in the interferometer prototypes, such as the 40 meter, the PNI or any of the tabletop setups. This lends some concreteness, but it should be clear that inappropriate analogies will be inevitable. An "operator interface" comprising discrete hardwired knobs, switches, oscilloscopes and test instruments has essentially unlimited signal bandwidth, but highly constrained display configuration and negligible automation. By contrast, the digital interface in LIGO will have arbitrary configuration flexibility and a great deal of automation, but may have restricted signal bandwidth (for a given channel or in aggregate). Operators and the controls themselves will have to adapt, and therefore a primary requirement (for both) is flexibility.

The following documents are referred to for background about the CDS system structure:

- LIGO-T950054-01-C, CDS Control and Monitoring Design Requirements Document
- LIGO-T950120-00-C, CDS Control and Monitoring Conceptual Design

2 LIVE DISPLAYS AND DIAGNOSTICS

Data display and diagnostic processes will organize information and actively probe the interferometer's operation to permit efficient assessment of the machine's state, characterize its parameters and internal degrees of freedom, and troubleshoot problems. Some of these functions are expected to be computationally intensive, so a flexible system is required which can trade off between update rate, process complexity, and active channel count and bandwidth.

2.1. Processed data presentation

2.1.1. Power spectra

General features:

- ~ 1024 points, DC-20, 200, 2kHz options
- "live" update, persistence, waterfall display options
- log mag vs. lin freq. (opt. log-log)
- "standard" axis range controls
- update procedure (incl. display update) completes in real time (i.e. t_proc < 1/2*f_max)
- cursors & markers
- trend plots of selected (marker) freq. bins or band averages to daughter display window
- zoom view to daughter window

- overlay “reference” spectrum trace (snapshot or recalled from library)
-

Process recipe:

  - digital LP filter time series, cutoff @ f_max
  - decimate/resample to get 2048 pts. in 1/2f_max stretch
  - vector multiply with Hamming window (raised cosine)
  - FFT-->1024 complex frequency points-take magnitudes
  - average each bin with last N FFT’s (N=0 to 100, say)
  - vector multiply with last swept sine calibration (for Strain signal)
  - do EU conversion (--> e.g. meters/root Hz for Strain signal)
  - update display

2.1.2. Pulse height analyzer display

The prototypes have not yet achieved the kind of “online” time domain pulse analysis capability which we agree is needed to characterize, understand and eliminate nongaussian events. I envision this online operator readout function as a stripped-down, streamlined version of the pulse search algorithm described in “Data Analysis Scenario and Data Flow Paradigm,” T960030.

General features:

- scans one or more channels for “significant” pulses (above selectable threshold)
- histograms number of pulses/energy bin vs. energy over previous N seconds for each chan.
- user selected prewhitening filter (from catalog)
- independently selected pulse template shape(s) (from catalog)
- simultaneous histograms for 1-5 template shapes
- configurable “dead-time” & multicount reject features (linked w/filter selections)
- can trigger O’scope display of buffered time series data (raw or filtered version) for “big” events above user-set threshold
- snapshot, store/recall, overlay reference plot

Process recipe:

  - hold unless data are qualified valid by IFO status bits (all $\overline{\text{LOCK}}$ low etc.)
  - filter selected channel time series through prewhitening filter
  - convolve filtered data stream with selected template(s)
  - for each template, square convolution output and compare w/threshold
  - if histogramming threshold exceeded, invoke deadtime/multicount reject algorithm to find “peak” value and time of event and report as single “pulse”
  - if “reporting” threshold(s) exceeded, flag oscilloscope process and pass corresponding waveform to it for display
  - update histogram(s) and refresh display

2.1.3. "Fast" Scope Displays:

General features:

Displays emulating a multichannel digital O'scope. Sensitivity, timebase, cursors, triggers, snapshot to daughter window/disk on trigger, freq. counter. Ability to hook into DAQ data stream channels or fast Remote Diagnostics digitizers (and configure digitizer input routing as required).

Inline filters (e.g. to improve visibility of high-frequency signals in LF noise or to improve triggering on fast, small glitches) will also be useful.

2.1.4. Time Series ("Trend") Displays:

General features:

- display converted to EU
- "now" (most recent sample) digital values displayed in plot corner (with trace legends for multiple traces)
- optional LP filters (selectable passband)
- "roll" (strip chart) principal mode; most recent point on right, old data drop off left
- typ. 500 pt. display resolution (hi/lo envelope or avg. if data pts. > display pts. for selected timebase)
- timebase ~ 1 sec/cm to 500 sec/cm on screen
- displays and traces individually selectable
- vertical scale selectable
- cursors

2.1.5. XY scope displays

General features:

Like an analog XY oscilloscope to display pairs of (slow) signals as dots in a plane. Primary examples are ASC alignment error signals. Useful features include:

- multiple signal pairs per display (coded by "dot" shape and color, text legend)
- offset adjusts
- gain adjust for each channel (track X&Y, independent X&Y option)
- "persistence" mode to give track of selected dot over recent past (period selectable)
- LP filters (selectable pass frequency) to clean up fuzz
- optional EU conversions
- "freeze" to hold dot positions on logic trigger or command

2.1.6. Video beam displays

General features:

- standard video rate (30 fps)
- moderate resolution adequate (e.g. 128 X 128)
- good bit depth/dynamic range (e.g. 16 bit/pixel)
- brightness, contrast, dark, gamma controls.

- pass TBD control signals back to video head (e.g. aperture, shutter, blanking, sync)
- metric/cursor capabilities useful, e.g.
  - cursor line slice to intensity vs. position graph in daughter window
  - centroid calculation & marker display
  - beam diameter calc
  - gaussian fit calc
  - intensity contour overlay or false color
  - mirror/flip

2.1.7. Audio output

The human ear/brain combination is highly developed at picking out weak signatures in noise and comparing weak tone frequencies. Listening to signals has proven to be a cheap and extremely powerful means of troubleshooting and debugging the prototypes. It should be fairly straightforward to send filtered versions of signals to workstation audio ports (for headphone monitoring or speaker broadcast).

General features:

- standard "built-in audio" characteristics (16 bit stereo @ 22 kS/sec/chan, say)
- selectable input source(s) (up to two for stereo comparison of signal pairs)
- selectable inline filter function (generally LP, HP, BP; simple "canned" shapes)
- optional AGC (adj. attack/decay) or selectable fixed volume
- optional blanking on trigger (e.g. cavity $\overline{\text{LOCK}}$)
- ability to "plug in" anywhere on the floor so you can listen while you smack things

2.2. Diagnostic stimulus/response test functions

Active stimulus/ response testing of interferometer systems will be a part of periodic standard calibration sequences, startup and alignment procedures, and most diagnostic investigations. For example, the primary calibration of the strain signal output will involve a swept sinusoidal test signal and/or a pseudorandom noise sequence applied to one TM suspension controller, a measurement of the response in the DAQ data stream, and signal processing to derive the complex transfer function (frequency response)[1]. In addition to DAQ system hardware, this type of test requires waveform generators near selected local units and analog multiplexers to direct their outputs to the correct stimulus points. For wideband responses, fast digitizers and associated input multiplexers will also be required. These additional degrees of freedom will be much easier to set up if graphical block diagram "test configuration" tools are available (see also 2.5).

Many of the desired functions are provided in dynamic signal analyzers like the HP3562/3 or SRS SR780. Emulating such dedicated signal analyzers may be a challenge, so it is probable that only a subset of capabilities can be provided. We may thus consider interfacing physical instruments so they can act as extensions of CDS. In any case, below is a summary of functional capabilities which will probably be used frequently.

---

[1] Calibration nominally belongs under the LSC's work scope, but this generic test style is used virtually everywhere.

2.2.1. Sinusoidal stimulus/power spectrum readout

Sine waves at selected frequencies and amplitudes are applied to one or more test points. The power spectra of output signals are monitored; by observing the relationship of the peak height to the noise background and knowing the size of the injected signals, one can immediately determine the SNR with respect to the injection point without separately measuring all the gain factors or interaction strengths (if the coupling mechanism is not well characterized or stable). Averaging of successive measurements is necessary to refine the power spectrum accuracy. Measurements of harmonics, two-tone intermodulation, and sideband amplitudes are also common.

2.2.2. Pseudorandom stimulus/power spectrum readout

This is functionally similar, but for nonlinear interactions[1] a sine wave test tends to underestimate the coupling. Using a noiselike test signal is often necessary to get the true interaction strength (i.e. that which applies for the "real" noise). A Gaussian random probability distribution is best for most purposes. Ability to bandlimit the random noise sharply (e.g. a "box" bandpass) will help to get around a serious problem: to get a measurable interaction, broadband excitation may introduce so much additional signal power that the system goes *really* nonlinear. Analyzers often can synchronize the pseudorandom sequence to the sampling and frame rates; this can improve the measurement dynamic range and variance by allowing coherent (vector) averaging of successive FFT measurements.

2.2.3. Pseudorandom stimulus/cross spectrum frequency response

This kind of test is principally meaningful for linear transfer functions. Stimulating a test point with a pseudorandom noise voltage, one collects two or more response waveforms and does a cross power spectrum for each pair. One of the response waveforms may be the stimulus itself, but frequency responses between any two affected points in the system may be needed. To insure that the cross power spectra measured are attributed to the same stimulus (and not naturally occurring noise on one input, for example) the quality of this kind of measurement needs to be monitored by averaging successive measurements and computing the "coherence function," a measure of how much power at each displayed frequency actually came from your stimulus. The magnitude and phase of the transfer function are read out on a Bode plot or similar. One will combine (multiply/divide) such transfer functions with power spectra from various sources, including ones with different numbers of points or nonuniform point spacing. Math operations on transfer functions and power spectra are also commonly used (like "divide by frequency squared" to turn force into displacement).

2.2.4. Swept sine frequency response

This method for measuring linear transfer functions just applies sine waves at successive frequencies[2] and computes the relative magnitude and phase of the response waveform. At each frequency the measurement may be repeated to reduce the effect of noise; one may also compute the coherence function to determine the quality of the result at each point, as above. It is obviously

---

[1] For example, coupling between mirror alignment and cavity length (quadratic) or amplifier saturation (limiting).

more time-intensive than FFT methods, but because the stimulus amplitude can be varied at each frequency, it can achieve much higher dynamic range (you can measure a transfer function that varies by 120 dB over the selected frequency range if you like, virtually impossible with FFT-based methods). Also, since only a narrow band is excited at any given time, this measurement usually adds negligible signal power to the system under test so it is much less perturbative. In addition to the cross spectrum capabilities mentioned above, swept signal analyzers generally include handy features like:

- active interpolation of additional frequency points when adjacent measurements differ markedly, so narrow features and transitions get accurately represented
- auto source level adjustment, to keep the signal at a manageable amplitude in one of the channels
- active dwell time, to be sure the measurement uses a reasonable number of cycles (or enough averages to get an acceptable variance) at each frequency

2.2.5. Triggered pulse response

This kind of measurement simply injects pulses of a desired shape into selected test points and triggers an oscilloscope readout to capture one or more responses (possibly including the stimulus itself). Pulse response waveforms give good indications of propagation delays and nonlinearities (especially rate limiting, which is hard to find otherwise). To an experienced operator it also gives a quick, intuitive summary of simple frequency responses.

2.3. Operator and data interfaces for dedicated instruments

Inevitably some jobs will need to employ special-purpose, dedicated instruments. At a minimum, things like RF and optical beam diagnostics will be done with direct laser beam or coax cable connections and special hardware that cannot practically be emulated by CDS. “Virtual” front panel operation and network transfer and storage of measurement data are required. While these functions are relatively routine nowadays, some “laboratory context” functions may have to be built, for example:

- recording the configuration of portable instrument I/O connections
- recording and associating interferometer state at time of a measurement with measurement data
- synchronizing instrument clocks and timestamping transferred data (time of capture vs. time of transfer)
- authority for changes in portable test equipment I/O configuration

2.4. Interferometer and subsystem “state-space” presentations

Status readout and control of state transitions can be facilitated by a user interface having a graphical representation of operation modes and transition paths. Highlighting the current mode and

---

2. Most useful if the amplitude for each frequency can be selected in advance to avoid overdriving resonances, for example; since the stimulus is also measured, this variation is divided out

"allowed" transitions (e.g. "cavity lock enabled" to "cavities locked" to "normal operation") will promote efficient startup and fault recovery. Explanatory depiction of interlock conditions *preventing* requested state transitions will also be helpful.

2.5. Sequence and procedure interfaces

Many calibration, measurement and diagnostic functions will require multistep sequences. These may involve hardware and software from several subsystems to work together properly. Automation can help accomplish these more efficiently, and will also insure consistent test conditions and parameters. Useful features may include:

- menu control of sequence parameters
- conditional branching based on data from prior steps
- automatic test scheduling at fixed times or on software cue (e.g. "recalibrate 2 minutes after each lock acquisition")

2.6. Alarm and error condition displays

Understanding alarms and error conditions quickly will require a clear, hierarchical interface. One desirable standard form is a graphical system block diagram; high level alarms may flash or change color of the icon representing the subsystem reporting the error. Lower level alarms may trigger opening of a window expanding the inner workings of the corresponding subsystem, to point directly to the problem.

This approach may be overwhelmed by "multiple alarm" scenarios, in which many systems report malfunctions due to a single common cause. Another layer of intelligence could be applied which "filters" alarms and points a diagnostic finger toward likely culprits. For example, an unusually large number of suspension controller signals hitting their limits might flag the system to display recent PEM seismometer records and check for dropouts on the crate power supplies. At the minimum, some crude inter-subsystem logic may be interposed to help reduce the diagnostic burden on the operator (for example, don't bother to report the power stabilizer signal is underrange if the laser is off!).

3 DISPLAY ORGANIZATION

To help assess the functional capabilities discussed above, it is useful to note how an interferometer operator might use displayed information to understand and troubleshoot the machine. The displays could be grouped by subsystem, e.g. PSL, SUS, ASC, LSC, plus perhaps PEM seismic or acoustic monitors etc. In this view, separate displays showing data from each of these subsystems appear to have equal status and at first glance the operator would need to cycle through all the components to understand the system's state. However, the implicit logic of the interferometer's' operation (Figure 1) lets an operator concentrate on "higher level" functions and dig deeper only where needed.

Expressing the logic hierarchy in the organization of subsystem displays will thus improve operator efficiency. An example of a "top level" display organization with a useful, but not too detailed, function subset is shown in Figure 2. Here preference is given to "higher level" functions the operator will want to start with as his/her point of reference. A further prioritization of the information within the subsystem displays (e.g. in hideable "layers" as done on the PSL Controls), as well as separate consoles focussed on individual subsystems, will augment this top-level condensed information and support offline test of isolated subsystems. This promotes efficient scanning and processing by a human operator without inhibiting deeper inquiry and diagnosis.

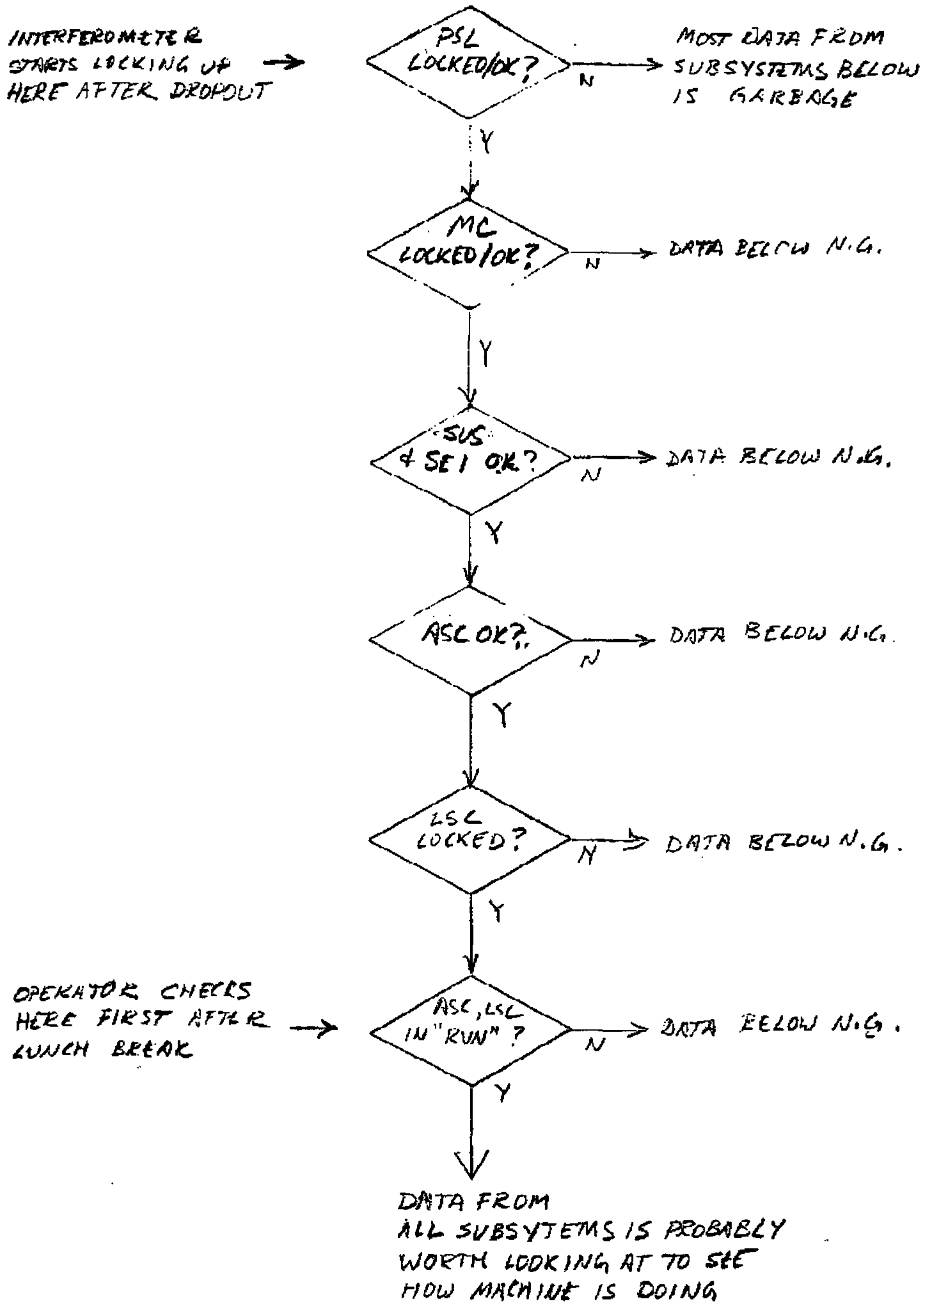

Figure 1: Implicit "troubleshooting" logic used by operator to decide what screens are worth looking at and in which order.

3/3/96 MEZ
3/4/96 MEZ

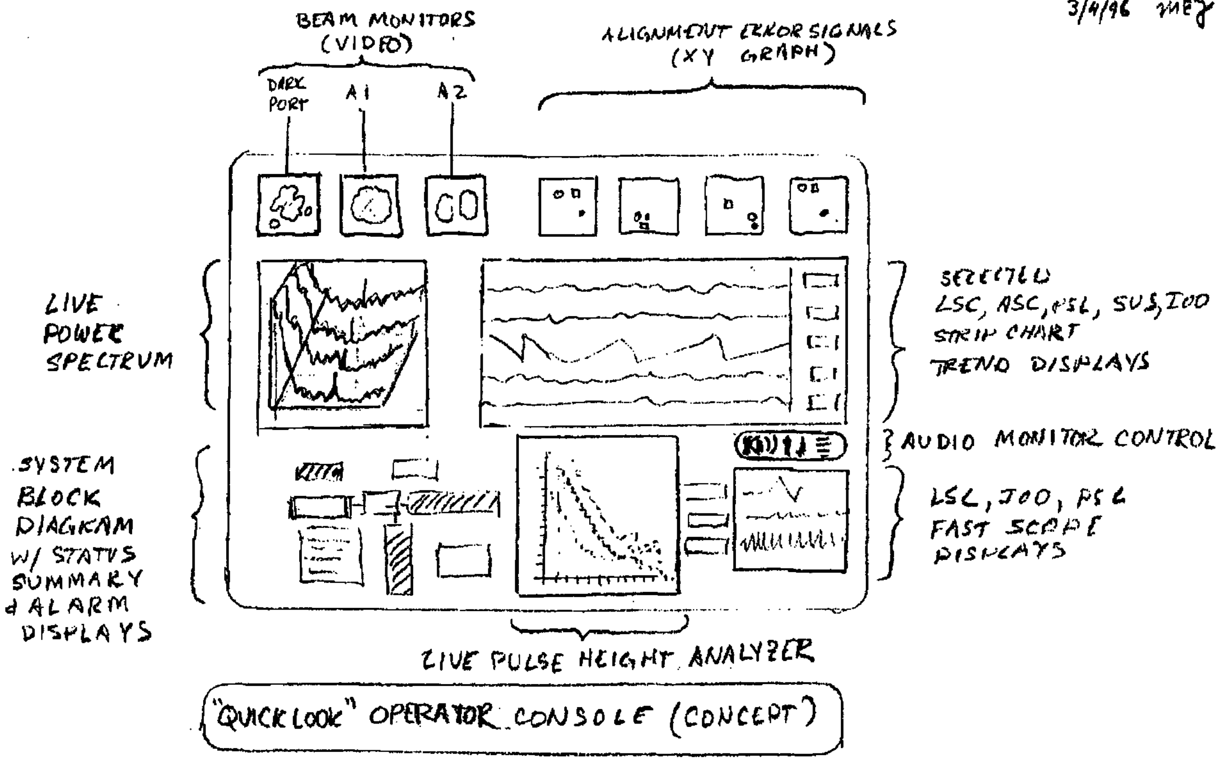

Figure 2: Top level realtime data display example