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

Technical Note	**LIGO-T000104-00 - D**	9/20/00

Monitoring Power Line Induced Artifacts at LIGO Hanford Observatory

Charles Shapiro

Distribution of this draft:

all

This is a technical 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/

Monitoring Power Line Induced Artifacts at LIGO Hanford Observatory

Charles Shapiro

Mentor: Dr. Daniel Sigg

LIGO Hanford Observatory

SURF 2000 Final Report

ABSTRACT

Power is one of LIGO's (Laser Interferometer Gravitational-wave Observatory) many sources of noise from the physical environment. Fluctuations and transients need to be monitored and logged so that their effects in the detector's data are not mistaken for astronomical events. As a solution, eleven voltage sensors were manufactured and installed at the power lines on various electronics racks throughout LHO (LIGO Hanford Observatory) so that voltage data could be digitized and made available online for computer analysis. A C++ program called MultiVolt was written to continuously download the data and make statistical calculations such as RMS, crest factor, frequency, and THD (Total Harmonic Distortion) for every signal. MultiVolt can run for weeks at a time as a background monitor – it saves each signal's results to an individual data file and posts the most recent summaries on a public web page. In the first week of observation, daily fluctuations were observed, and the trends were consistent with the hypothesis that higher power usage by local residents late in the day (for air conditioning, cooking, lights, etc.) effects small variations in the regional power.

Introduction

LIGO uses very sensitive equipment in its attempt to detect gravity waves – so sensitive that local disturbances can easily cause noise to leak into what it sees. One source of noise comes from the power lines that supply LIGO Hanford Observatory (LHO) with its electricity. When fluctuations occur in the power lines, they have to be recorded and logged so that their effects in the detector are not mistaken for astronomical events. One goal of this project was to build Physical Environment Monitor (PEM) voltage monitors that would allow the power in each building to be measured by an analog-to-digital converter (ADC) and broadcast over various channels. A second goal of this project was to write a C++ program that would continuously monitor the channels and give periodic summaries of the power quality. Both goals were achieved – the power lines at 11 locations throughout LHO are currently being monitored by "MultiVolt", the final program.

This text will begin with a discussion of power quality and the various statistical quantities measured by MultiVolt. Then the PEM voltage monitors and the process of data collection will be described. Next, there will be a description of MultiVolt's fundamental structure and operations. Finally, the results of the program will be discussed.

Power Quality

Since LIGO is coming closer to stabilizing its detectors and the output of its lasers, power quality is becoming an increasing concern. Although "power" typically refers to a rate of energy transfer, in this context "power" is essentially synonymous with "voltage". The power company delivers LHO (as it does most places) a sinusoidal voltage with an amplitude of about 170 volts and a frequency of about 60 Hz. The stability of this voltage is essential to the stability of the machines that use it – when harsh changes occur in the voltage, *power quality problems* can develop. A power quality problem can be defined broadly as a "power problem manifested in voltage, current, or frequency deviations that results in failure or misoperation of … equipment." (DMB, 2) A LIGO "misoperation" that is cause for concern is the distortion of data due to transients or large fluctuations in the power.

Power quality problems can arise from a variety of sources, the most obvious of which is the power company. For instance, one of its generators can fail, causing LHO to experience a voltage sag for a few moments. Other sources are natural phenomena like lightning strikes, which can cause voltage spikes. Power quality problems can also be caused in the buildings where the power is used – a faulty transformer can distort the waveform, and a short circuit can cause voltage sags. Yet another source of problems is the surrounding community – the ways that other people use their power (e.g. say everyone within 30 miles turns on their heaters on a cold day) can affect how power arrives at LHO.

The power lines have a number of quantitative properties that MultiVolt can measure to check for power quality problems. A basic one is the *root mean square* (RMS) of the voltage, which is given as

$$V_{RMS} = \sqrt{\frac{1}{T}\int_0^T V(t)^2 dt} \qquad \text{Equation 1}$$

where *V* is the voltage as a function of time and *T* is some interval of time containing several periods of V. The RMS gives the effective DC voltage level of the AC voltage. Unlike the average of the AC voltage, the RMS voltage is not zero in general – it is usually about equal to the amplitude divided by the square root of two, which is the RMS of a perfect sine wave. Obviously, a high RMS means that there is a lot of power, and a low RMS means that there is not. A common problem to watch for is an unexpected drop in RMS.

Another basic property measured by MultiVolt is *crest factor*. Crest factor is simply the absolute value of the maximum voltage (usually the amplitude) divided by the RMS:

$$\mathrm{CF} = \left|\frac{V_{\max}}{V_{RMS}}\right| \qquad \text{Equation 2}$$

The crest factor is usually about equal to the square root of two – large deviations from this are an indication that the waveform does not look like a perfect sine wave. In addition, a sudden change in crest factor without a noticeable change in RMS can indicate a transient.

MultiVolt also determines the frequency of the voltage, which does not stay constant at 60 Hz. The frequency might be determined by merely counting the number of crests in an interval and dividing by the interval length, but this method has several drawbacks: it assumes an integral number of wavelengths is being used; and it will readily misinterpret distortions that contain multiple crests. A more efficient way would be to consider the Fourier Transform of a few seconds of the voltage. In the neighborhood of 60 Hz, it will look like Figure 1 below.

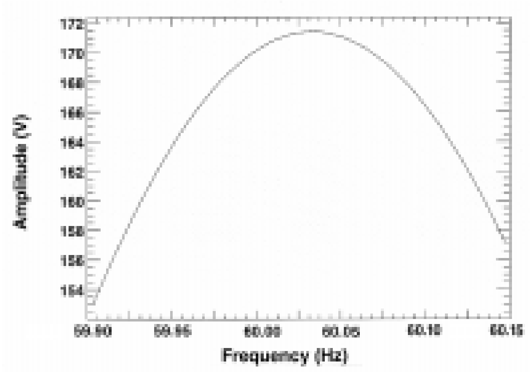

Figure 1 – Fourier Transform of Five Seconds of Voltage

The shape of the region in Figure 1 is approximately a parabola. If three points on the parabola in the neighborhood of 60 Hz are determined, one can then solve for the equation of the parabola and then for the location of the maximum – this is precisely what MultiVolt does. It finds the complex Fourier coefficients of three frequencies using

$$C(f) = \frac{2}{T}\int_{0}^{T} V(t) * e^{-i2\pi ft}\,dt \qquad \text{Equation 3}$$

where *f* is the frequency and *T* is the length of the recorded interval. The equation of the parabola is of the form

$$\mathrm{a}f^{2} + \mathrm{b}f + \mathrm{c} = \left|C(f)\right| \qquad \text{Equation 4}$$

The coefficients a, b and c can be found by plugging *f* and *C*(*f*) into Equation 4 for all three frequencies and using a bit of linear algebra. The frequency with the largest amplitude is then given by

$$f_{\max} = -\frac{\mathrm{b}}{2\mathrm{a}} \qquad \text{Equation 5}$$

When measuring frequency, large fluctuations are not expected since the frequency is directly determined by the motion of massive generators owned by the power company – it is not likely that the generators will accelerate or decelerate suddenly. However, it is important that the frequency be known so that it can be used in the calculation of harmonics.

Harmonics are frequencies that are an integer multiple of the fundamental frequency (which is of course 60 Hz for power lines). When the shape of the 60 Hz wave has a constantly repeating distortion, the amplitudes of its harmonics will be substantial in a Fourier Series of the entire voltage. Below is a sample of the imperfect LHO voltage and its corresponding frequency spectrum – in this case, only the odd harmonics are affected.

Figure 2 – The Waveform of LHO's Power and its Frequency Spectrum

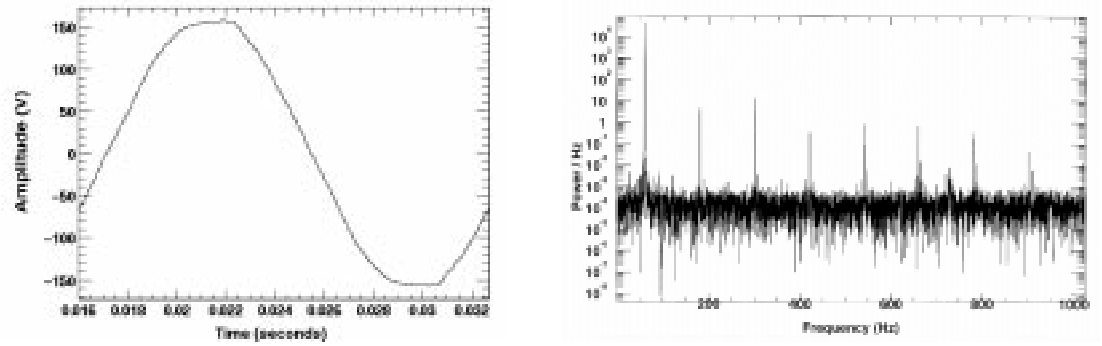

MultiVolt applies a Hanning window to the data to ensure that signal lines to not "leak" into neighboring bins. The equation of the window is

$$w(t) = -\frac{1}{2}\cos(t\pi / T) + \frac{1}{2} \qquad \text{Equation 6}$$

The program then uses Equation 3 to calculate the amplitudes of the harmonics (as percentages of the fundamental frequency's) and their phases (relative to the fundamental frequency). The amplitudes are

$$A_h = \left| \frac{C(f_h)}{C(f_{fund})} \right| \qquad \text{Equation 7}$$

and the phases are

$$\phi_h = \mathrm{Arg}[C(f_h)] - \mathrm{Arg}[C(f_{fund})] \qquad \text{Equation 8}$$

Since the amplitudes in Equation 7 are relative, no renormalization for the application of the Hanning window is required. Nevertheless, the amplitudes are actually equal to $|C(f)| * \sqrt{1.5}$ - this is taken into account when the amplitude of the fundamental frequency is reported. MultiVolt also calculates the *total harmonic distortion*, a measure of the harmonics' effective magnitude. Total harmonic distortion (THD) is expressed as a percentage:

$$\mathrm{THD} = \frac{\sqrt{\sum_{h=2}^{h_{\max}} V_h^2}}{V_1} \qquad \text{Equation 9}$$

where V_1 is the amplitude of the fundamental frequency and V_h is the amplitude of the h^{th} harmonic. Harmonic distortion is often due to the non-linear characteristics of devices and loads on the power system.

(DMB, 24) Examining the harmonics and the THD is a good way to troubleshoot mechanical errors. A THD less than 5% is usually considered acceptable. The THD of the waveform in Figure 2 is 3%.

PEM Voltage Monitors

The PEM voltage monitors, which later became known as “sensors”, are devices used to allow the voltage to be measured by ADCs. Power is delivered to the buildings at LHO in three distinct phases, each separated from the other two by 120 degrees. At the corner, mid, and end stations, sensitive electronic equipment on various electronics racks are supplied with instrumental power that needs to be monitored. The power source cannot be *directly* linked to the ADC: the typical amplitude of AC power line voltage is about 170 volts, and the ADCs can only accept ±2 volts. Thus, only a small proportion of the voltage must be passed to the ADCs. The sensors are designed to do just that – they plug directly into the power lines and transmit about one one-thousandth of the voltage to the ADCs. Each sensor can link two voltage phases to an ADC. A diagram of the LHO sensor locations and the phases they monitor is shown in Figure 3 below.

Figure 3 – Sensor Locations

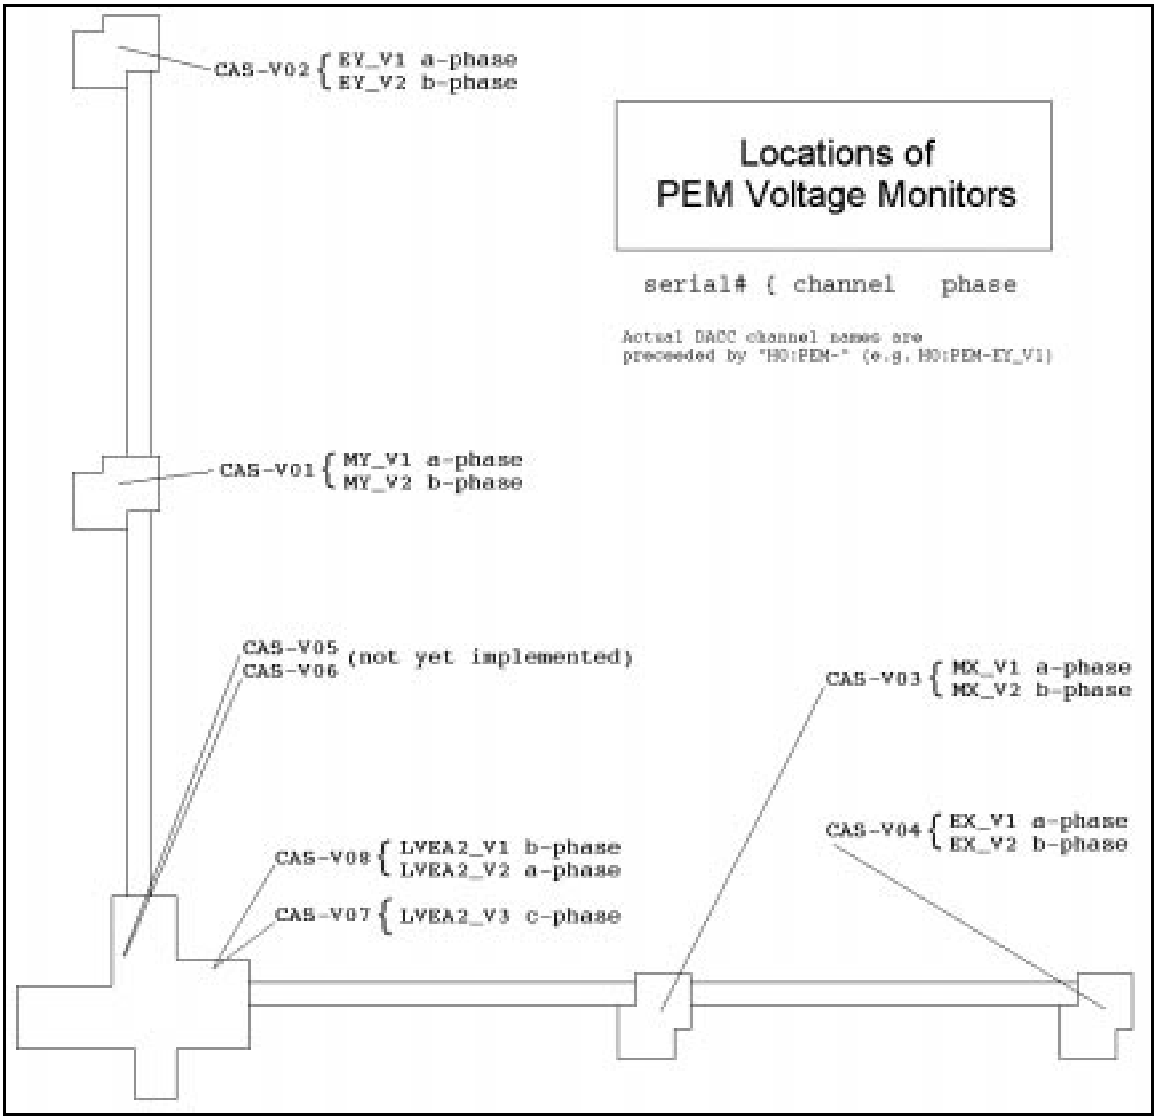

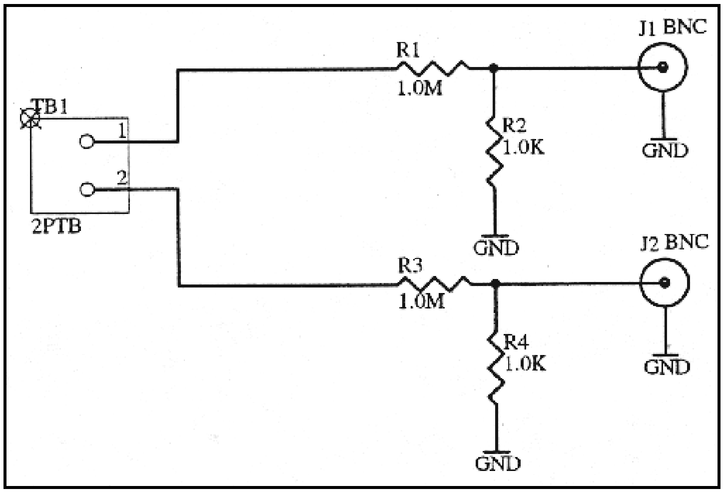

Figure 4 – Basic Diagram of a Sensor

Figure 4 shows a basic sensor diagram (a picture of a sensor can be found in Appendix A). The terminal block on the left is attached to a power line – path 1 is connected to “hot” while path 2 is connected to “neutral”. The British Naval Connectors (BNCs) are linked to an ADC that will measure the voltage difference between them. For the most part, the voltage in the “hot” path 1 will fluctuate relative to the “neutral” path 2, whose voltage will remain constant. To see how the sensor reduces the measured voltage, we start by realizing that almost none of the current in path 1 reaches the BNC, which is linked to an ADC with very high impedance. We then have the approximation

$$V_0 = I_0(R_1 + R_2) \qquad \text{Equation 10}$$

where V_0 and I_0 are the total voltage and current in path 1. The voltage drop at the BNC must equal the voltage drop at R_2 , thus

$$V_{BNC} = V_2 = I_2 R_2 = I_0 R_2 \qquad \text{Equation 11}$$

where V_2 and I_2 are the voltage drop and current at R_2. Then, plugging in V_{BNC} / R_2 for I_0 in Equation 10 and solving for V_{BNC}, we get

$$V_{BNC} = \frac{V_0 R_2}{R_1 + R_2} \qquad \text{Equation 12}$$

Since R_1 and R_2 are 1.0 MΩ and 1.0 kΩ respectively, the voltage difference measured by the ADC will be about one one-thousandth of the actual voltage difference in the power line (i.e. about ±170 mV). Each sensor actually contains two arrangements like the one shown in Figure 4 (equipped with fuses) so that it can monitor the power at two adjacent electronics racks.

With the sensors installed, the power can be measured and digitally recorded. Each sensor’s output is first sent through an anti-aliasing chassis, which applies an 850 Hz low pass filter (the filters are 8-pole elliptical analog filters). The filtered data is then sampled and timestamped at 16.384 kHz by ADCs, which are synchronized by the Global Positioning System (GPS). The ADCs send their data to data collection units (DCUs) that organize the data into separate channels – each monitored power line gets its own channel of data. The DCUs also decimate the data by eight, making the final sampling frequency 2048 Hz. Every DCU transmits its data to the central data acquisition (DAQ) controller, which prepares the data to be broadcasted by channel to various systems. One of these systems is the Data Monitoring

Tool (DMT), which reads the data into shared memory where it can be downloaded by MultiVolt and other software programs.

Sensor Calibration

After the sensors were manufactured and installed but before they could be used, they needed to be calibrated. The outputs of the ADCs are reported in ADC units, and for each sensor, the ratio of ADC units to volts had to be established. In addition, each sensor's DC offset needed to be determined – zero volts does not correspond to zero ADC units. Since the relationship between ADC units and volts is linear:

$$ADC\,units \;=\; ratio \,*\, volts \;+\; DC \qquad \text{Equation 13}$$

the calibration is quite straightforward. The DC offset is simply the ADC output when the voltage is zero – it can be found by terminating each sensor's input and recording the corresponding ADC values using software on LHO's Global Diagnostics System. The ratio can then be found by supplying some known voltage and again recording the ADC output – doing this gives two points on the line expressed by Equation 13. The sensors were connected to a constant 100.0 volts using a Hewlett Packard E3612A DC power supply, giving

$$ratio \;=\; \frac{ADC\,units\;(\text{volts}=100.0)\;\text{-}\;ADC\,units\;(\text{volts}=0)}{100.0\;volts\;\text{-}\;0\;volts} \qquad \text{Equation 14}$$

A calibration table (taken from an LHO engineering log) is shown below in Figure 5. It is organized by channel name (i.e. by power line – see Figure 3 above for a diagram of power line and sensor locations).

Figure 5 – Calibration Table for LHO Voltage Sensors

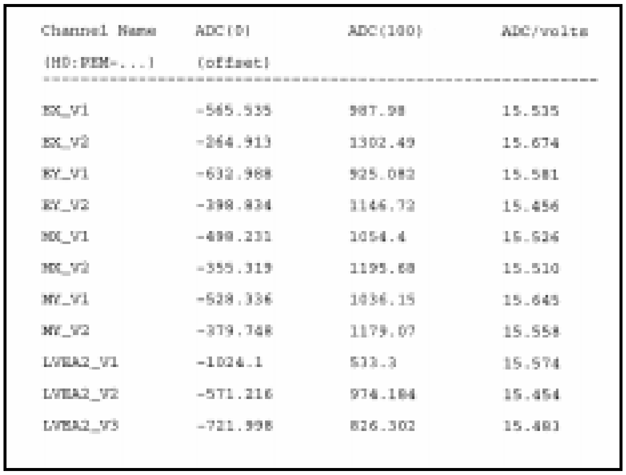

Channel Name (H0:PEM-...)	ADC(0) (offset)	ADC(100)	ADC/volts
EX_V1	-565.535	987.98	15.535
EX_V2	-264.913	1302.49	15.674
EY_V1	-632.988	925.082	15.581
EY_V2	-398.834	1146.72	15.456
MX_V1	-498.231	1054.4	16.526
MX_V2	-355.319	1195.68	15.510
MY_V1	-528.336	1036.15	15.645
MY_V2	-379.748	1179.07	15.558
LVEA2_V1	-1024.1	533.3	15.574
LVEA2_V2	-571.216	974.184	15.454
LVEA2_V3	-721.998	826.302	15.483

The MultiVolt Program and its Results

With the data available on channels in the DMT, the MultiVolt program can download the information continuously and make its calculations. Its default settings are to look at 10-second segments of data and average its results over a minute – this is done simultaneously for all channels. The statistical quantities calculated are those mentioned in the Power Quality section above: RMS, crest factor, frequency, the amplitudes and phases of harmonics, and THD. The harmonics are only found up to the Nyquist frequency; since the final sampling frequency is 2048 Hz (see PEM Voltage Monitors), the Nyquist frequency is 1024 Hz, thus 17 harmonics are calculated. Where the computation of these quantities requires an integral sum (Equation 1 and Equation 3), the program uses simple numerical versions of the calculus equations (trapezoid rule integration). MultiVolt's results are saved to data files every minute – each channel has it's own data file. The most recent results are written to a public web page each minute. The web page is a useful tool for quickly diagnosing power quality problems at LHO, and long-tern trends can be investigated in the data files by loading them into Matlab™. For more information about the MultiVolt program, see Appendix B.

The decision to use 10-second segments of data for calculations was not arbitrary: MultiVolt's early results revealed that longer segments of data yielded unrealistic results for the harmonics. The problem was traced back to the frequency of LHO's power: the frequency wanders so much that the *average* frequency, which is used to compute the harmonics, carries very little weight in the frequency spectrum. This causes the amplitudes of the harmonics to appear lower than they are. Below is an example of MultiVolt's results.

Figure 6 – Frequency vs. Time for the Y-Mid Station (a-phase) Starting on 7/30/00

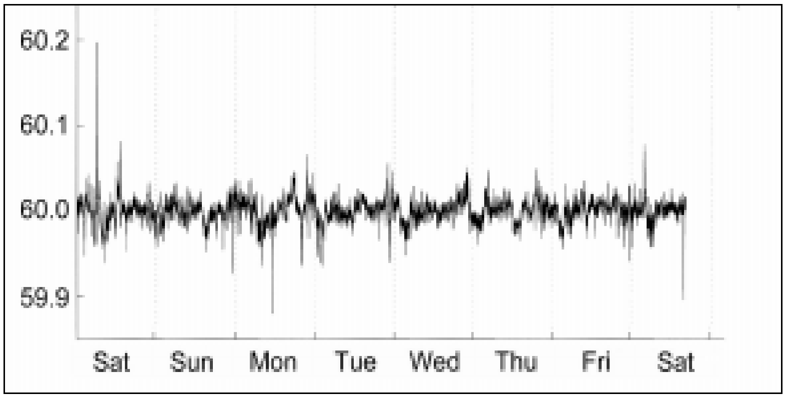

Although Figure 6 was plotted using data from the Y-mid station a-phase, the plot for every other power line is nearly identical. It is clear from the graph that the frequency is very active although it only varies by about 0.3% from the mean. It was originally uncertain whether these tiny fluctuations were real or whether they were the results of computational or data collection errors. To take a closer look at the character of the frequency, the average frequency and amplitude of a long voltage time series was found. Then, a perfect sine wave with this frequency and amplitude was subtracted from the original data so that the beat envelope, shown in Figure 7, could be observed. It appears solid black because the waves are spaced extremely close together.

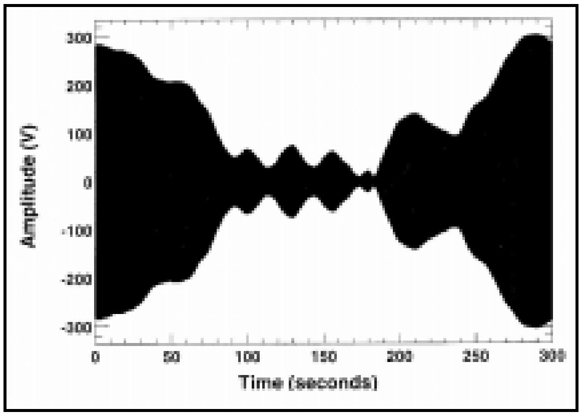

Figure 7 – The Envelope Formed When a Perfect Sine Wave is Subtracted from an LHO Voltage Time Series

This barely resembles the smooth envelope formed by subtracting two waves with similar, constant frequencies. The original wave's frequency is clearly changing, causing an erratic phase shift between the two waves that creates an undulating shape. The widths of the undulations are approximately 30 seconds; this indicates that 30 seconds is too long to wait to take a measurement – the frequency wanders too much. Thus, the measurements should be kept to 10 seconds or less.

MultiVolt was run on the DMT for a week, and the results were analyzed examined for long-term trends. Some of the trends are shown below in Figure 8. The scales of the figures are arbitrary – each plot was translated and scaled so that it could be viewed comprehensively with the others. The divisions along the horizontal axis are made at the midnights.

Figure 8 – Statistical Trends in LHO Power Lines for One Week, Starting on 7/30/00

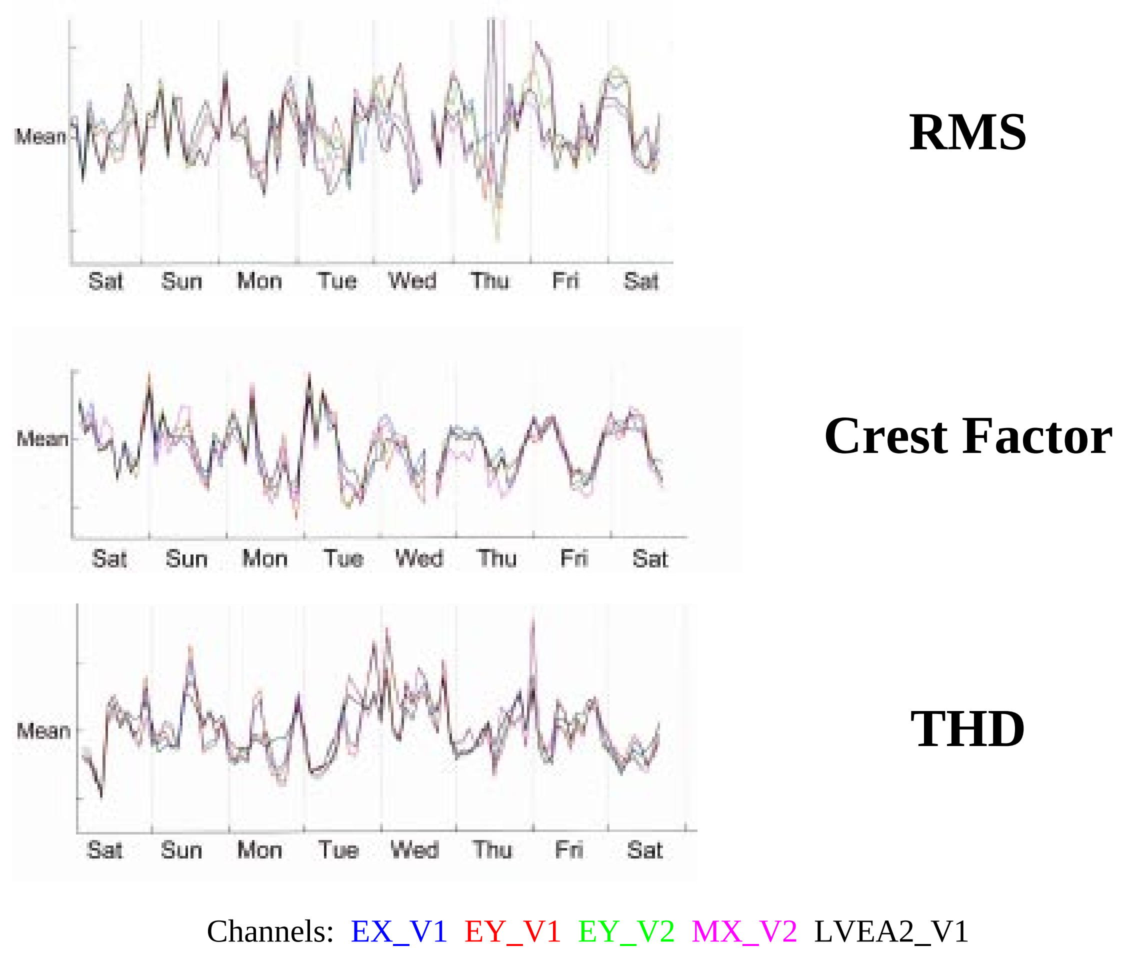

Although not indicated above, the mean RMS ranged from about 117 V in the LVEA to 121 V in the mid and end stations – variations from the mean were generally less than 2%. The average THD was usually about 3% in the LVEA and 2.5% in the mid and end stations with variations less than 15%. The mean crest factor was approximately 1.4 everywhere. A daily pattern is somewhat apparent in the RMS: it seems to be lower in the evening than it is in the morning. The crest factor correlates very well with this trend – the THD does so to a lesser degree. A plausible explanation is that local Hanford residents effect these changes by using more power later in the day. Cooking appliances and especially air conditioners are generally used more in the afternoon – this puts a strain on the power company, which may not be able to keep the voltage exactly constant as the electricity usage increases.

Notice in Figure 8 that there was an RMS spike on Thursday. This corresponds to a real overvoltage that occurred on certain b-phase power lines. The RMS abruptly rose about 3 volts for several hours and then returned to normal. That is precisely the type of power quality problem that LHO will want to identify with MultiVolt. Had the overvoltage occurred during an actual run of the LIGO detector, it most likely would have rendered some of the gravitational wave data unusable.

Conclusions

The PEM voltage monitors have been installed and are currently collecting data from 11 locations throughout LHO. The MultiVolt program was used to produce a periodic list of statistical quantities for each power line, and it was running for days and weeks at a time on the DMT background. The output of the program is useful to diagnose power quality problems as demonstrated by the first week of running at LHO. The next step of this project would be to add a "trigger" to the program. Once the normal limits of the RMS, frequency, THD, etc. for each channel are determined, the program can check to see whether these limits are ever exceeded. If they are, it can save the abnormal data to a separate file and issue an alert to the control room. The program can also be adapted to monitor current - it could then find quantities such as "power factor" using both the current and voltage data. PEM current monitors would of course have to be built and installed first.

Acknowledgements

I would like to thank my mentor, Dr. Daniel Sigg, for *everything* (the list is too long). Thanks also to Dave Barker and John Zweizig for keeping the DMT hardware and software running smoothly. Special thanks goes to Josh Meyers and Richard McCarthy for all of their help with the manufacturing and installation of the sensors, and to Johnathan Oberg for handling some of the "dirty work."

This work is supported by the National Science Foundation cooperative agreement PHY-9210038

APPENDIX A – A PEM Voltage Monitor

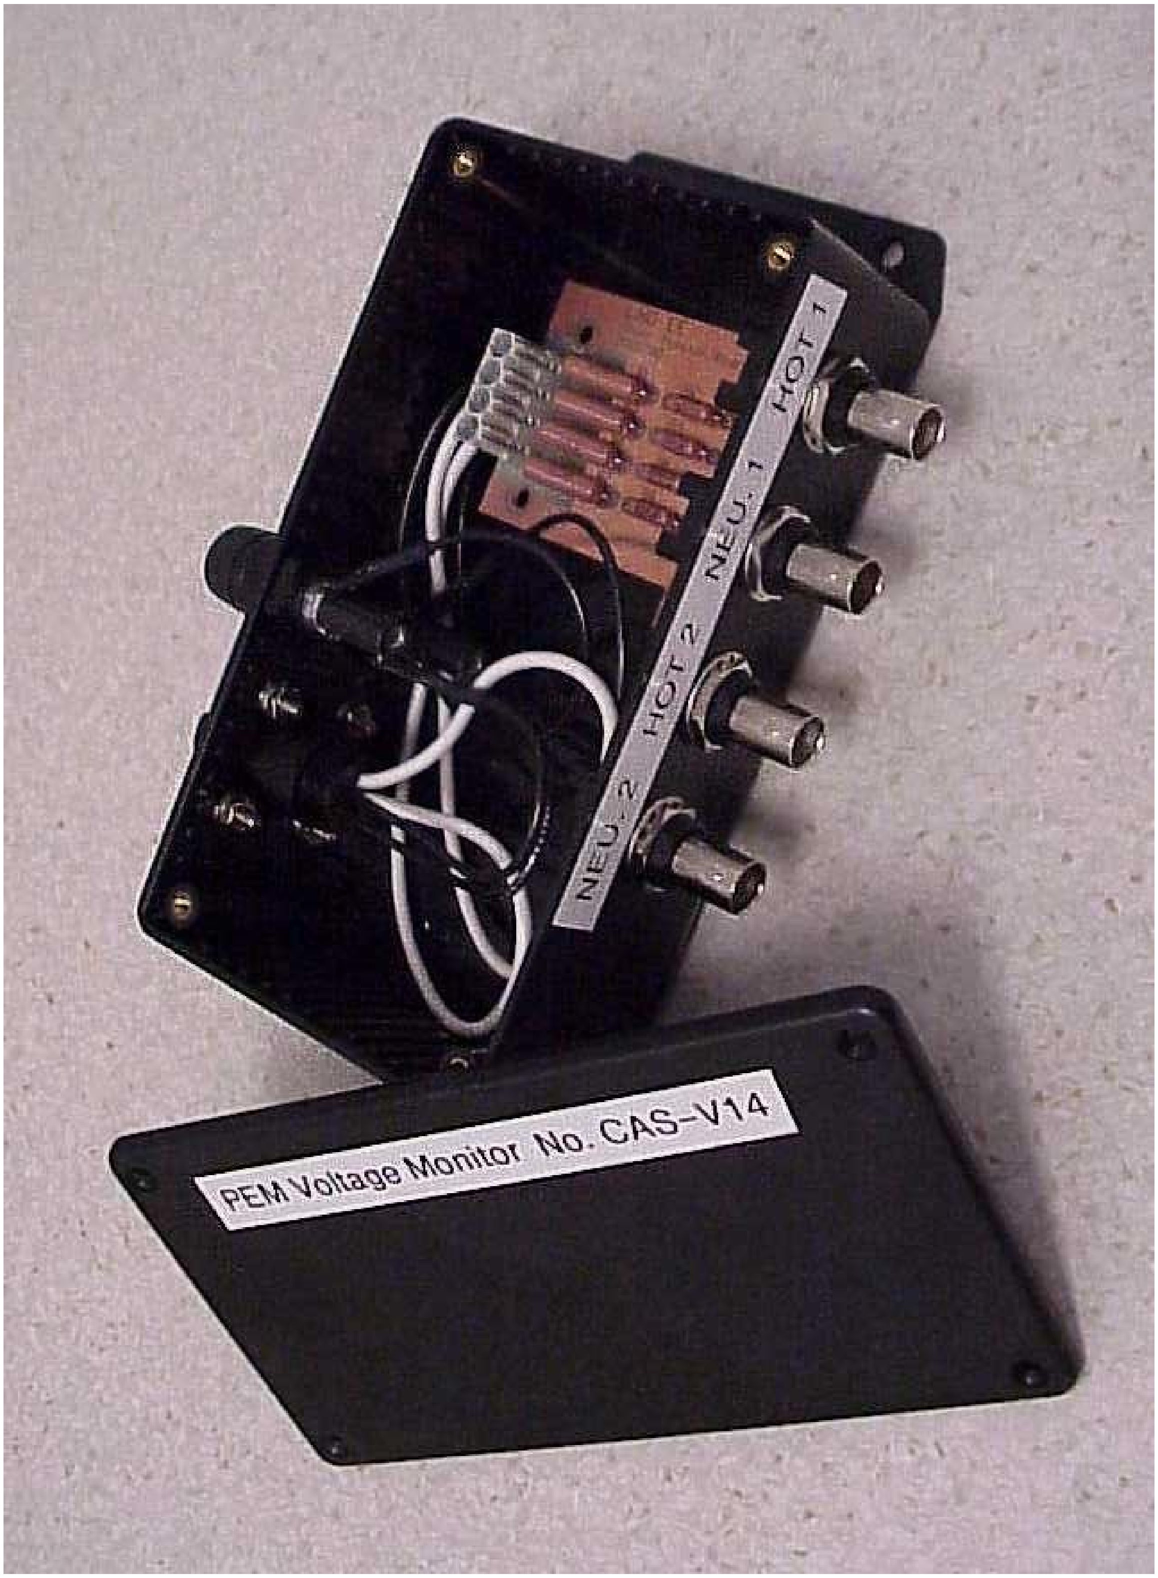

APPENDIX B – The MultiVolt Program Code Structure

The MultiVolt program started as a collection of ROOT macros and evolved into an automated C++ program. Its ability to monitor and store data for several data channels at once required an object-oriented approach. A class called VoltWatcher was designed to download voltage time series from a channel and calculate all of the quantities described above. A class called VoltWriter was designed to average the results over a minute and write them to the screen, a data file, or a web page. The main program, MultiVolt, controls an array of VoltWatchers (one for each channel) and uses a VoltWriter to append individual data files (one for each channel) and update a web page with the latest results.

MultiVolt does not merely download large segments of data and perform all of its calculations on every channel at once. It makes calculations each second for one or two channels, thus evenly distributing its tasks. To accomplish this, the VoltWatcher class required a system of buffers. Each VoltWatcher has two 10-second buffers – a "passive" buffer that is filled with 10 seconds of data and an "active" buffer that is only partially filled. Every second, all of the active buffers are filled with one second of data. Then calculations can be quickly made on data in the passive buffers for one or two channels. When the active buffer is filled with 10 seconds of data, the passive buffer is cleared and the roles are switched (active becomes passive and vice versa).

The program codes and documentation can be viewed here:

http://`www.ligo-wa.caltech.edu/gds/multivolt`

References

Dugan, Roger C., McGranaghan, Mark F., Beaty, H. Wayne. Electrical Power Systems Quality. McGraw Hill, 1996. (DMB)

Press, William H., Teukolsky, Saul A., Vetterling, William T., Flannery, Brian P. Numerical Recipes in C. Cambridge Univ. Press, 1993.

Schildt, Herbert. C++ The Complete Reference. Osborne McGraw Hill, 1995.