Detecting grid instabilities using real time synchronized Phasor calculations[1]

C. H. Wells[2]

Abstract

This paper demonstrates how real time phasor calculations can be used to detect imminent grid instabilities. This information can be used to prevent blackouts, by advising grid operators to increase grid damping by controlling reactive power. Data transmission latencies are automatically corrected before calculations are performed. This is especially important in computing synchronized relative phase angles. Fast Fourier Transforms are performed in real time and used to identify increasing modal amplitudes. Alarms are automatically raised when peak amplitudes of these modes violate statistical quality control rules. An example is shown using real data from a recent blackout.

Introduction

There is an increasing interest in using synchronized phasor measurements to aid in controlling electric power grids. Phasor measurements are available from a new class of instrumentation called Phasor Measurement Units (PMU). These devices measure absolute phase angles according to IEEE-1344 standards and include both current and voltage in each phase, in addition to a wide range of other measurements. The Arbiter 1133a is one such device considered to be a PMU: it measures 684 variables in total, at rates up to 240 Hz. The most commonly used variables are voltage, current, real and reactive power, frequency, and phase angle.

Each measurement is time tagged to an accuracy of better than 1 microsecond using the top of second signals from the GPS. The phase angle measurements are absolute relative to this signal. This allows information from different substations to be used for Wide Area Measurement applications. The phasor measurements are reported as complex numbers consisting of a real and imaginary component. Some devices convert this information to polar form, per recommendations of IEEE-37.118. Regardless of the form of the phasor data, difference calculations are required to determine information about the grid.

The basic equations describing the voltage and current in a single phase are[3]

$e(t) = \sqrt{2}E(t)\cos\left[\omega t + \delta(t)\right]$ and

---

[1] Presented at: Second Carnegie Mellon Conference in Electric Power Systems, “Monitoring, Sensing, Software and Its Valuation for the Changing Electric Power Industry, January 11-12, 2006, Pittsburg, PA.

[2] Platform Architect, OSIsoft, Inc. San Leandro, CA.

[3] Notation used is from Ilic and Zaborszky, “Dynamics and Control of Large Electric Power Systems, John Wiley and Sons, Wiley Interscience, New York, 2000.

$$i(t) = \sqrt{2} I(t) \cos(\left[\omega t + \delta(t) + \varphi(t)\right],$$

where $\sqrt{2}E$ and $\sqrt{2}I$ represent the peak values of voltage and current, ω is frequency in radians per second, $\delta(t)$ is the reference angle in radians, and $\varphi(t)$ is the power factor angle. In the real world, these change with time, and hence the quasi-stationary. PMU's measure these variables faster than 20 times per second.

With these measurement rates, one can determine transient stability and generator and area swings that are generally slower since load and generator inertias that are generally slower than (1/20) of the speed of the carrier frequency.

The instantaneous power is defined as

$$P(t) = e(t)i(t) = p(t) + q(t) = E(t)I(t)\cos\varphi(t)\{1 + \cos 2\left[\omega t + \delta(t)\right]\} + E(t)I(t)\sin\varphi(t)\{1 + 2\left[\omega t + \delta(t)\right]\}$$

where $p(t)$ and $q(t)$ are the real and reactive components of the total power $P(t)$. PMUs measure both real and reactive power as well as power factors, phase angles, and rms values of current and voltages on each phase.

Each PMU is required to provide a UTC time stamp for each measured variable. UTC time differs from both GPS time and from Atomic time. Two times each year, UTC time may change by one "leap" second. For example, at the end of 2005, one leap second will be added to UTC time. Thus far all leap seconds have been positive.

UTC-TAI = -33 seconds.

where TAI is International Atomic Time. Changes, if required, are made simultaneously so, in effect, UTC time is smooth.

Thus all data from the PMUs are supplied with UTC time and hence can be used in WAM systems to perform time synchronized calculations.

Real time calculations: Latency issues

PMUs are typically located in substations and transmit data via packet switched networks to a central server. There is significant time delay between the time the packet is sent and the time it arrives at the server. Typically this varies between 15 ms and 200 ms.

In order to compute the real time relative phase angle between two substations, both the minuend and subtrahend must have the same time stamp. Thus

$$\delta_r(t) = \delta_m(t) - \delta_s(t)$$

where the m and s subscripts represent the minuend and subtrahend and the r subscript represents the relative phase angle. The time (t) must be exactly at the same time stamp.

One property of the synchronized PMU time stamps, is that they are relative to the top of second mark. In most GPS receivers, this pulse is accurate to ±200 ns. This means that time stamps from disparate PMUs can be aligned exactly to a time accuracy of better than 1 μs. An example of a raw data dump of PMU data into an Excel spreadsheet is shown below.

Date Time (seq)	Ch A V	Phase(deg)
12-04-03 16:47:27.050 (01)	107.425	180.509
12-04-03 16:47:27.100 (02)	107.41	180.224
12-04-03 16:47:27.150 (03)	107.409	179.958
12-04-03 16:47:27.200 (04)	107.411	179.685
12-04-03 16:47:27.250 (05)	107.415	179.418
12-04-03 16:47:27.300 (06)	107.41	179.151
12-04-03 16:47:27.350 (07)	107.407	178.878
12-04-03 16:47:27.400 (08)	107.418	178.618
12-04-03 16:47:27.450 (09)	107.42	178.345
12-04-03 16:47:27.500 (10)	107.413	178.072
12-04-03 16:47:27.550 (11)	107.421	177.793
12-04-03 16:47:27.600 (12)	107.418	177.526
12-04-03 16:47:27.650 (13)	107.418	177.265
12-04-03 16:47:27.700 (14)	107.409	176.992
12-04-03 16:47:27.750 (15)	107.404	176.713
12-04-03 16:47:27.800 (16)	107.422	176.452
12-04-03 16:47:27.850 (17)	107.419	176.191
12-04-03 16:47:27.900 (18)	107.408	175.923
12-04-03 16:47:27.950 (19)	107.412	175.65
12-04-03 16:47:28.000 (00)	107.416	175.388
12-04-03 16:47:28.050 (01)	107.425	175.109
12-04-03 16:47:28.100 (02)	107.414	174.829
12-04-03 16:47:28.150 (03)	107.418	174.549
12-04-03 16:47:28.200 (04)	107.431	174.288
12-04-03 16:47:28.250 (05)	107.426	174.019
12-04-03 16:47:28.300 (06)	107.414	173.738
12-04-03 16:47:28.350 (07)	107.419	173.44
12-04-03 16:47:28.400 (08)	107.417	173.14
12-04-03 16:47:28.450 (09)	107.416	172.847
12-04-03 16:47:28.500 (10)	107.414	172.566
12-04-03 16:47:28.550 (11)	107.407	172.272
12-04-03 16:47:28.600 (12)	107.402	171.984
12-04-03 16:47:28.650 (13)	107.41	171.697

Table 1. Raw phasor data

Note also that PMUs also send a sequence number associated with the time stamp. This is a number from 0 to 19 representing the time offset from the top of second pulse, with 0 representing the measurement at the top of second.

Computing historical differences is straight forward since no time interpolation is required: simply request data from the archive at the exact time stamp values, modulo 50 ms. However, computing the differences in real time is more complex.

But, it is impossible to compute PMU differences in real time. However, the calculations can be performed as soon as both the minuend and the subtrahend arrive at the host computer. The differencing algorithm must buffer both the minuend and subtrahend values and then perform the calculation as soon as the oldest argument arrives. This can be done by examining the time stamp values as they arrive at the host computer.

Using the public Internet, the oldest value is typically no older than about 150 ms; however, this varies greatly and hence cannot be assumed to be a constant value. A dedicated network will improve the worse case latency; however, it will still be no better than about 25 ms and will also be non-deterministic. Since the latency is generally less than one second, the sequence numbers could also be use to align the data. However the data can be buffered by more than one second when there problems on the network; hence, we chose to use to the actual time stamp value to perform time comparisons.

The latency issues address above are handled automatically by the PI-ACE software (Advanced Computing Engine, a component of the OSIsoft PI server).

Real time calculations: Angle discontinuity

PMUs report phasor measurements as either or both polar or rectangular (complex) coordinate systems. Both are discontinuous in the angle measurement since its range is (0-360) degrees. This presents difficulties when performing subtractions. Typically the absolute angle measurements vary slowly either increasing or decreasing at rates of -5 to + 5 degrees per second; thus, every 2 or 3 minutes there will be a discontinuity; i.e., the absolute phase angle will jump from 0 to 360 or vice versa. This condition is also handled in the PI-ACE software. An example of the discontinuity is shown below

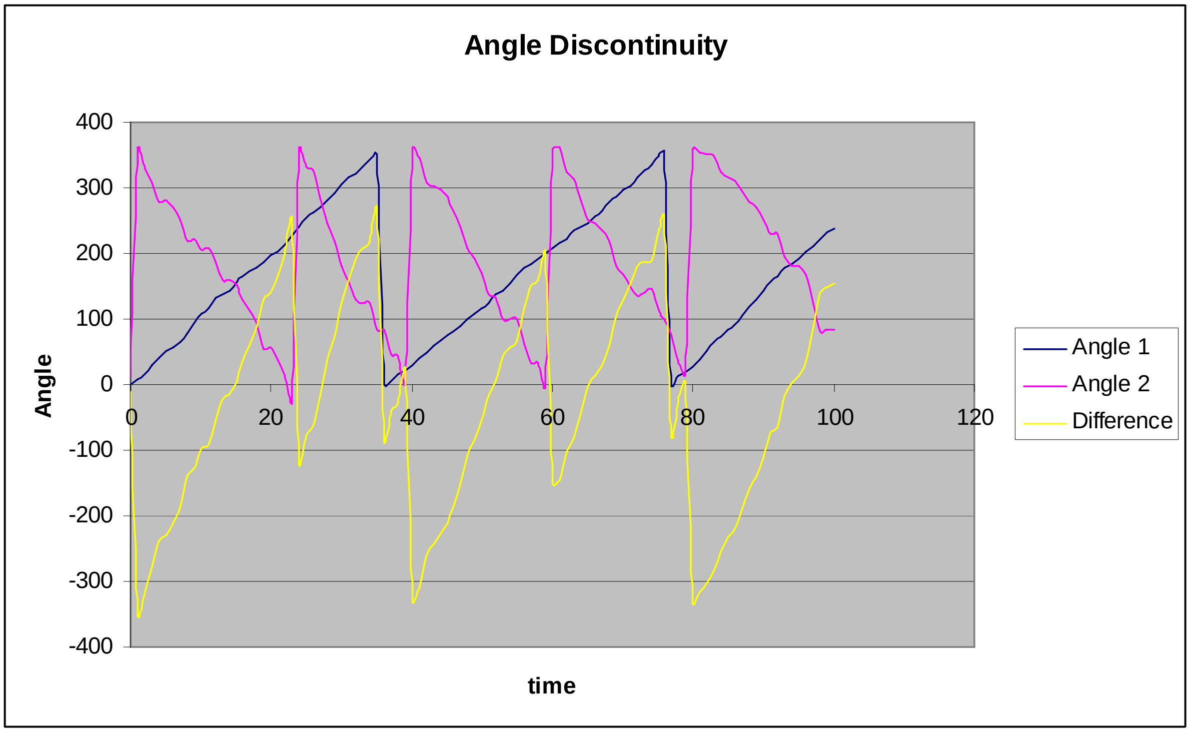

Figure 1 Angle discontinuity

An example of typical variations in phase angle is shown in the top trend Figure 2. This shows frequency and positive sequence voltage angle for a period of about one hour. Notice that the direction of rotation of the phasor changes during this one hour period. This is common in all of the system we have examined. An increasing absolute phase angle corresponds to a net frequency greater than the nominal and correspondingly, a negative slope corresponds to a "droop" below the nominal frequency.

The second set of traces represents the three voltages from the 480 volt building feeder. These voltages are all in excess of the nominal 277.1 volt on each phase. The voltage drops correspond to the current spikes on the third set of traces on this display.

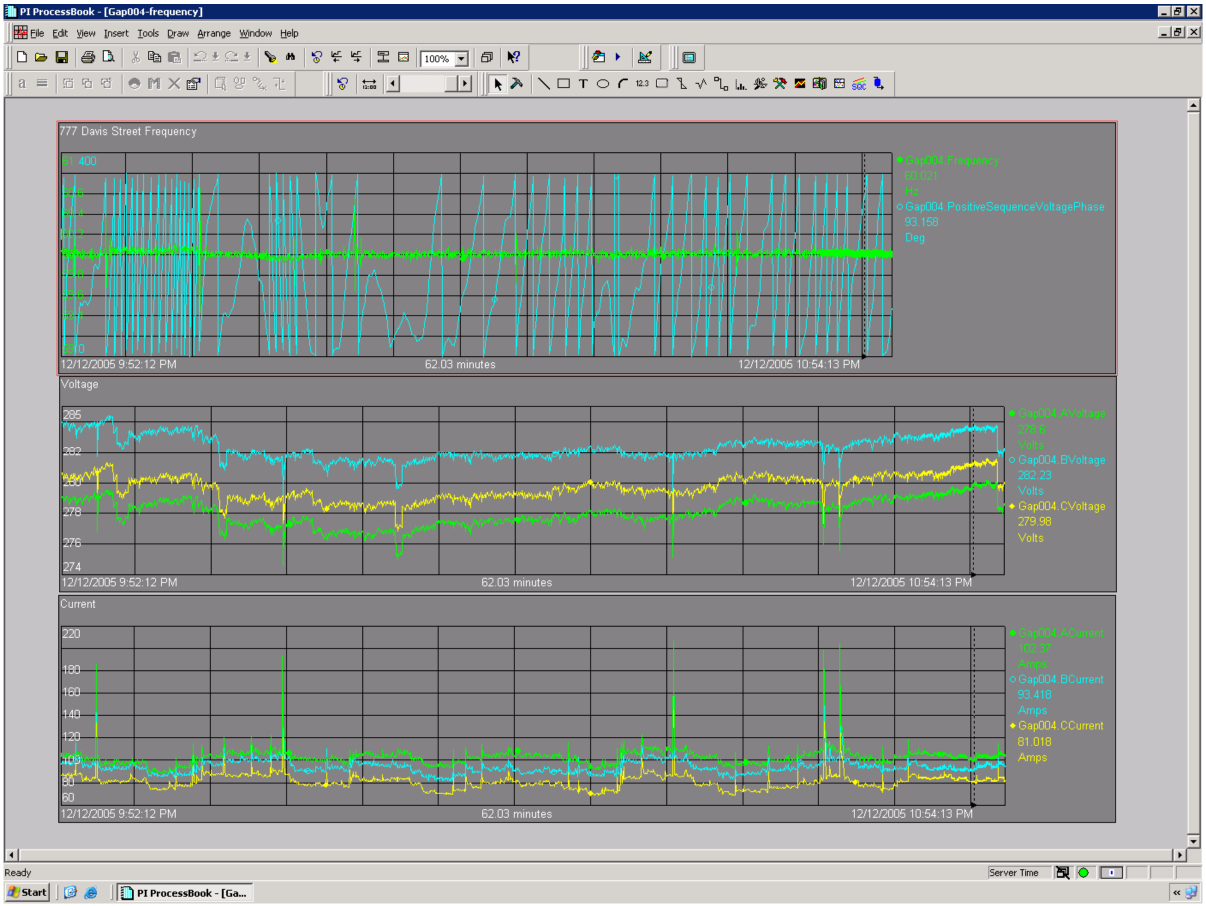

Figure 2 Typical phase angle behavior in a 3 phase feeder circuit.

Real time calculations: Fast Fourier Transforms

The electric power grid is a complex system, consisting of thousands of stochastic loads and generators. It is a system that constantly oscillates at known frequencies, for example, in the Eastern Interconnection common oscillations occur at 0.3 and 0.7 Hz. These are caused by the both the circuit topology and the character of the loads and sources.

Power system interconnections are generally over damped with damping ratios greater than one at most modes. However, when the system impedances are such that the inductive and capacitive reactances are nearly equal at a given frequency, the system can become unstable. This is often called a Hopf bifurcation.

One method of determining stability in any dynamic system is to analyze the periodic behavior of measured state variables, frequency being one convenient measure. Moving window Fast Fourier Transforms (FFT) provide a convenient means of determining stability in real time. This type of analysis is equivalent to wavelet analysis, since time variation is taken into account by the moving window.

Suppose we are computing FFTs in near real time at 20 Hz rates. Then if the peak amplitude in any mode is continually increasing, the system damping is decreasing and the system is nearing instability. Additionally, if the area under the FFT is increasing, the total noise level in the system is increasing and the system is becoming less stable.

Using frequency as the key measured variable, the following phenomena are well known from control theory: (1) the area under the FFT curve should be nearly zero if the system is “in control”, this means that the frequency measurements behave as if they were a random Gaussian process with a mean of zero, (2) any growth of a peak in the FFT spectrum should not grow without bounds, (3) any peak in the FFT spectrum should also exhibit Gaussian behavior.

There are several properties to consider when computing FFTs on line: (1) evenly spaced data, (2) number of gaps in the data, (3) width of the gap, and (4) total percentage of missing data. The window width is also important since it sets the range of modes that can be extracted from the moving window. Normally, we set: (1) the maximum number of gaps to 10, (2) the largest gap not exceeding 15 measurement points, and (3) no more than a total of 10 percent missing data in each window.

In our studies, we normally use 4096 points in the moving window with even spacing of 50 ms providing a moving window of 204 seconds. The FFT calculation over this window computes modal amplitude from 0.005 Hz to 10 Hz.

A sample of typical frequency data used in the FFT calculation is shown in Figures 3 and 4 below. Figure 3 shows frequency waveforms over a five minute interval.

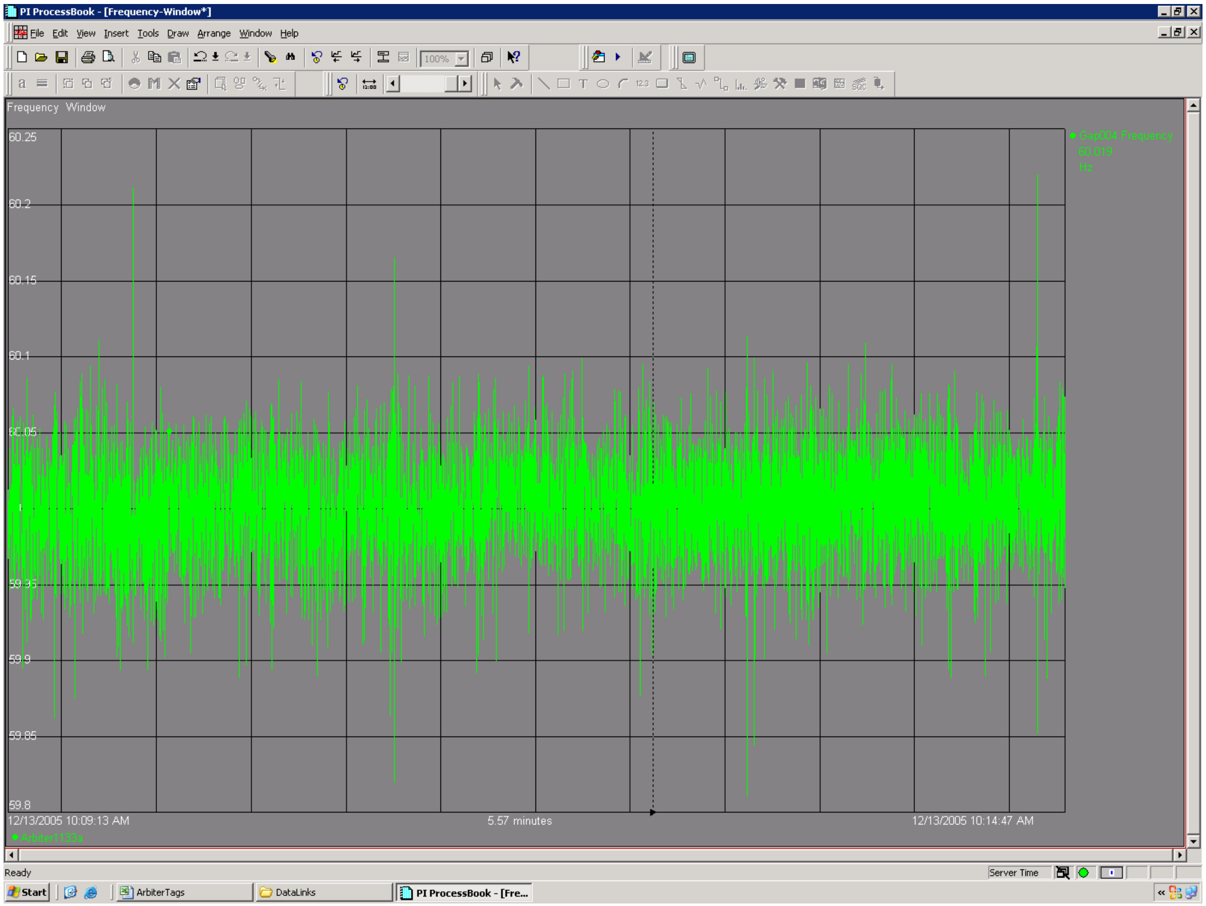

Figure 3 Approximately two windows of frequency data.

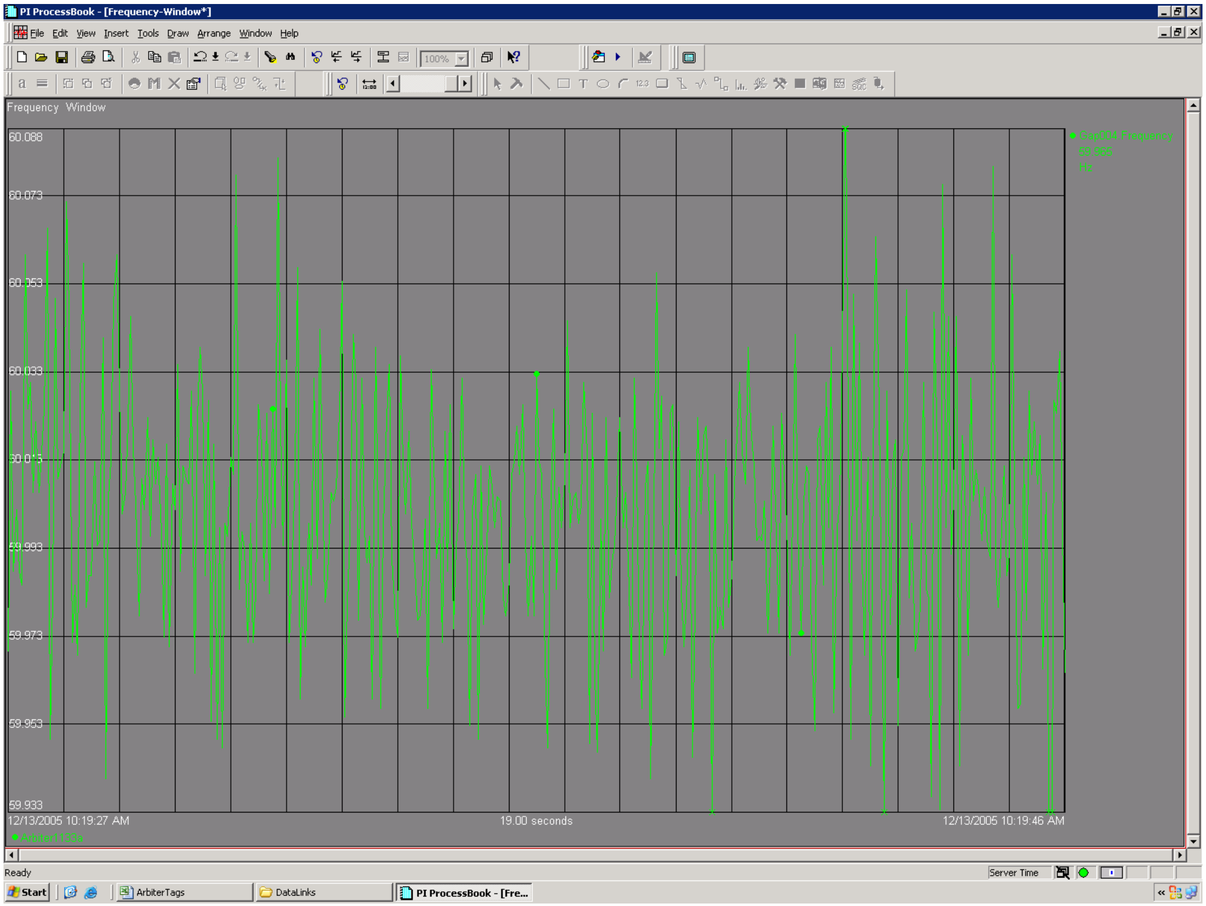

Figure 4 Zoomed in view of frequency

We compute the peak height, peak location, and peak width of the largest N peaks in the spectrum as part of the FFT algorithm: typically N=10. The resulting data are stored as PI tags so a complete history of the spectrum can be characterized.

A sample of a waterfall chart illustrating the FFT over time is shown in Figure 5. This shows how the spectrum changes over time (time range is about 41 minutes in this example). Note the location of the three major peaks. These were precursors to a blackout that occurred 25 minutes after the occurrence of the first peak.

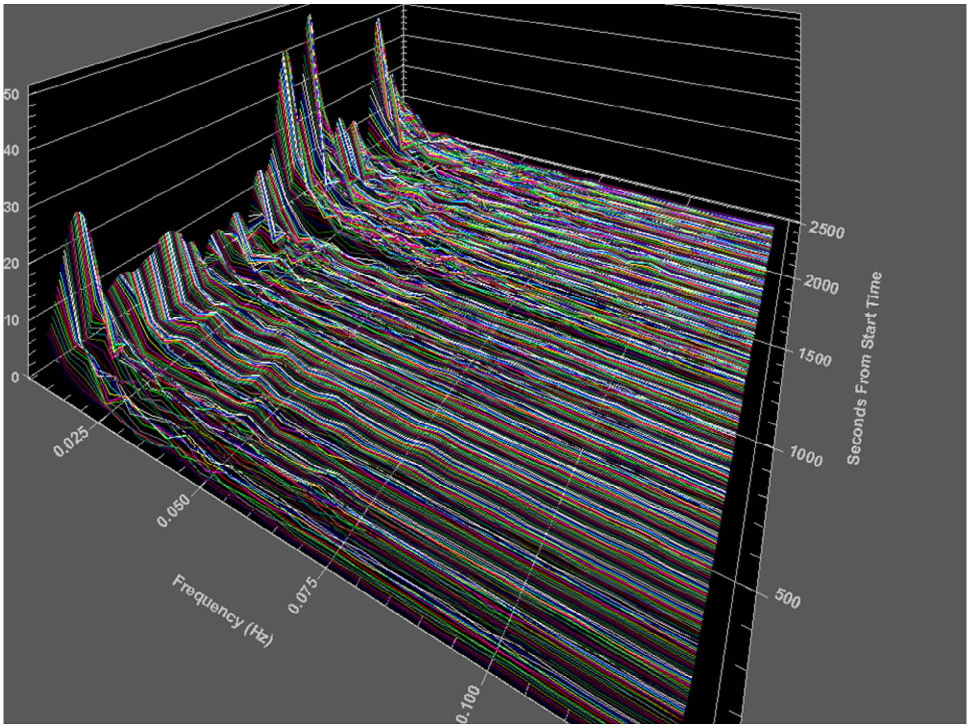

Figure 5. Real time waterfall chart

The top 3 peaks are shown in Figure 6. Each peak corresponds to a “mode” or Eigenvalue in the grid.

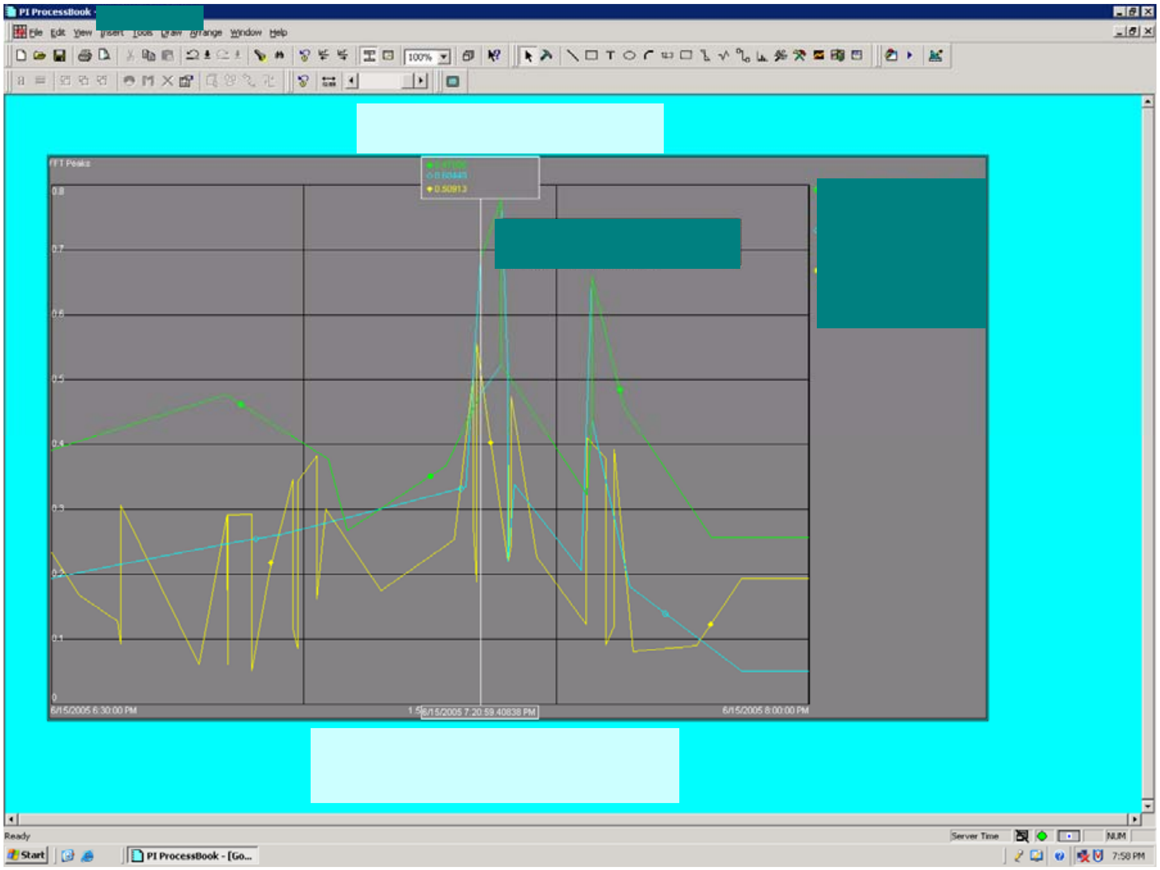

Figure 6a. Largest three peaks in the FFT spectrum over time

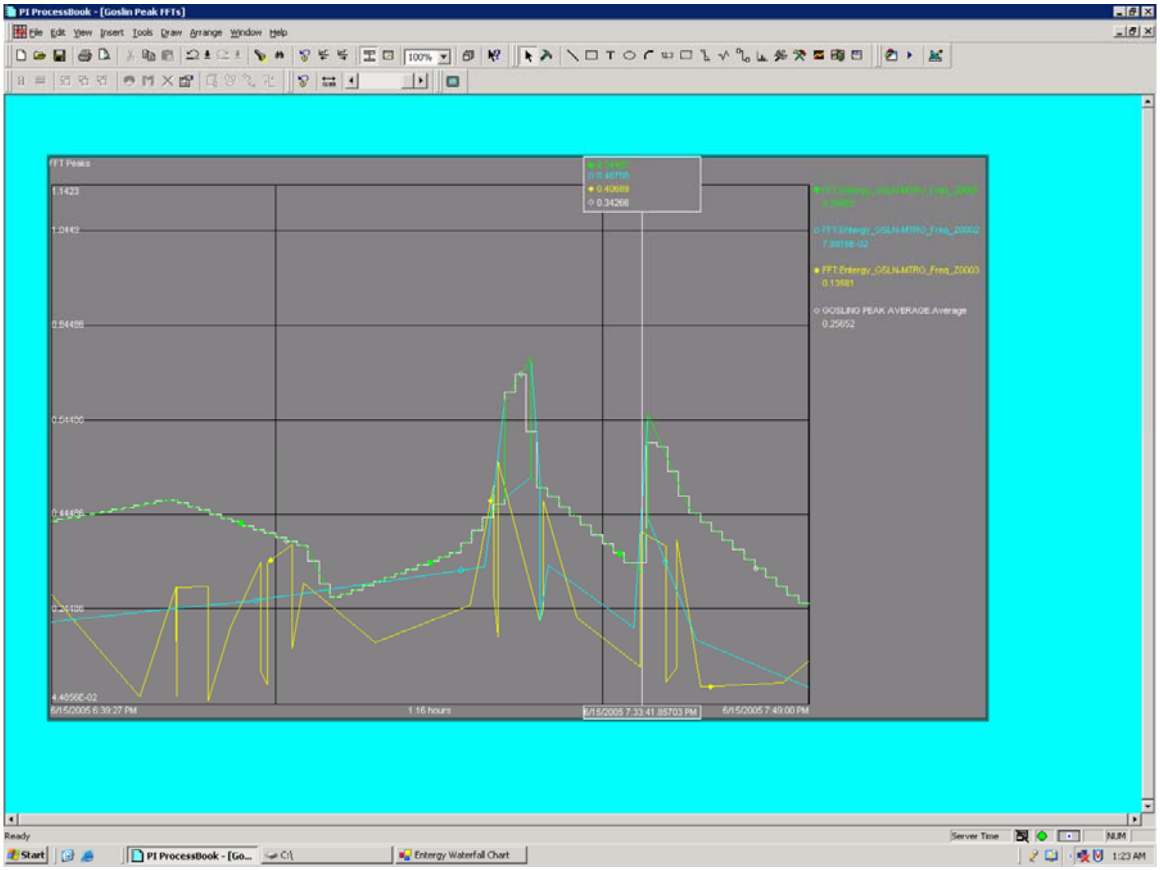

Figure 6b. Largest three peaks with dynamic averaging on first peak.

Damping calculations:

A classical non-linear analysis of power systems begins with a model defined as

$$\dot{x} = f(x, y, p)$$
$$0 = g(x, y, p)$$

where x is a vector of the dynamic state variables, y is a vector of the algebraic state variables, and p is a vector of the system parameters. The f and g functions are non-linear functions of the states and parameters. This system is called an algebraically constrained ordinary differential equation describing the dynamics of power system. The eigenvalues of this system can be determined around a linearized operating point. The eigenvalues move in the complex plane (similar to root-locus) as the system topology and parameters vary. The system becomes unstable when any real part of any eigenvalue in the system appears in the right half plane. A bifurcation occurs when the eigenvalues are purely imaginary. This defines a Hopf[4] bifurcation point.

---

[4] Zaborszky, J., et. al, "Computing the Hopf Bifurcation boundary", EPRI Report, RP3573-10, 1995.

We assume there are N fundamental frequencies in the waveform each represented as a complex eigenvalue. Rather that attempting to solve for the eigenvalues, assuming a valid model existed, we extract the fundamental frequencies then look at their time varying characteristics. Clearly if the damping ratio decreases, the FFT associated with that peak increases implying that the system is becoming unstable.

One could use time domain approaches to determine damping. This requires, first determining the order of the system, and then estimating parameters continually on moving data windows. This does not seem to yield useful results since the order of the system changes with changes in grid topology.

The damping coefficient at each harmonic can be determined by examining the history of the spectrum. If a peak increases in amplitude over time, the grid damping is negative, i.e. the system is going unstable; i.e., the poles of this mode are moving into the right half plane. We suggest computing the derivatives automatically using moving window polynomial filters. This method computes first and second time derivatives in the noisy environment.

Hopf bifurcations[5] are discussed in detail in Ilic's text, Chapter 9. Ilic used PEALS, and PSSLF4 to confirm a Hopf Bifurcation caused the Rush Island blackout of 1993 (in this case the topology changed, an insulator at a substation failed). Hopf bifurcations are basically when there are no real parts of the system eigenvalues: this is the bifurcation limit. Any movement to the right will cause the system to collapse.

Since, we do not know the actual grid topology at any moment in real time, or the non-linear characteristics of the components in the model, nor do we know the stochastic loads and generation, we think the best approach to grid stability is to extract this information from accurate measurements from PMUs. We also demonstrated how to use cross correlations between geometrically separate frequency measurements to determine grid coherence. The next level of stability could include computation of the spectral content of relative phasors.

We use moving window polynomial filters to compute the first and second derivatives of the peaks: this is required since the spectra are noisy. The polynomial window width is the "persistence time:" it can be set to any value, for example 5 or 10 minutes, etc. Additionally, since the polynomial coefficients are computed in the filter, we can use these to "predict the future." This is very similar to load forecasting. We forecast the peaks 5 to 10 minutes in the future. For example, suppose the 108 harmonic peak rate is increasing at 2 times the rate of the adjacent peaks? This indicates potential collapse. The thresholds for limits can be adaptive as well, i.e. they continually ratchet upward. Also for this to work, all other NERC standards must be satisfied; i.e., voltage, current, and frequency have to be within limits.

Calculation of the damping coefficient for each major peak

---

[5] Abed and Varaiya, in "Int. J. Electric Power Energy Systems, Vol 6, pp 37-43, 1984

Classical second order systems are represented by the equation:

$$\ddot{x} + 2\zeta\omega_0\dot{x} + \omega_0^2 x = 0$$

where ζ is the damping coefficient and ω_0 is the undamped natural frequency. If the damping coefficient is greater than one, the system is said to be over-damped, if less than one but greater than zero, the system is under-damped, if equal to zero, the system is oscillatory, and if less than zero, unstable.

However, the grid is more complex than a single second order system with a single natural frequency. Hence we need to compute the damping coefficient for each natural frequency in the system. This can be done using the FFTs as outlined above.

The system uses the history of the spectrum to compute the time derivative of the natural frequencies in the spectrum.

Let $\phi_i(k)$ = the magnitude of the spectrum at harmonic number i at time (k). Let $\dot{\phi}_i(k)$ be the first derivative with respect to time of this value. Then

If $\dot{\phi}_i(k)$ < 0, system is stable

If $\dot{\phi}_i(k)$ = 0 system is near instability, a bifurcation

If $\dot{\phi}_i(k)$ > system is unstable

These conditions have to be persistent to trigger an alarm. The persistence time is a configurable number, typically far less than the NERC 30 minute persistence limits.

The peak amplitudes vary over time and contain noise. Hence the computation of the time derivatives is done using a moving polynomial filter, typically second order. This produces both the first and the second derivatives and is also used to predict the value of the peak at short periods ahead, typically up to 30 seconds. This is done using the polynomial coefficients and future time values. We commonly use the linear polynomial as the prediction function, rather than second or higher orders.

The damping coefficient at each fundamental harmonic can also be computed direct from the FFT since the FFT and the Laplace transform are directly related by the square function. So by squaring each element of the FFT, we get the value of the Laplace transform.

The damping coefficient is related to location of the poles of the system in units of the Laplace transform as follows[6].

$$\alpha = \frac{IM}{RE} = \frac{\sqrt{(1-\zeta^2)}}{\zeta}$$

where IM and RE are the real and imaginary parts of the poles of a second order system as represented in the S plane.

Or solving for ζ

$$\zeta = \sqrt{\frac{1}{(\alpha^2+1)}}$$

Thus, the damping coefficients can be computed in real-time from the moving window LaPlace transforms. This can be done for each harmonic number of significance

Phase Portraits:

Using the X-Y plot functions within ProcessBook, phase portraits can be generated in real time. These are plots of frequency versus relative phase angle. Examples are shown in the Figures 7a and 7b. These are plots of relative phase angle on the x axis and frequency on the y axis. The time range of interest is selectable and in these plots is about 5 minutes. The most recent reading is shown in yellow and the next to the most is shown in orange. Thus the implicit time history is shown in the chart.

---

[6]B. Bequette, Process Control: Modeling, Design and Simulation, Prentice Hall, Dec 2002.

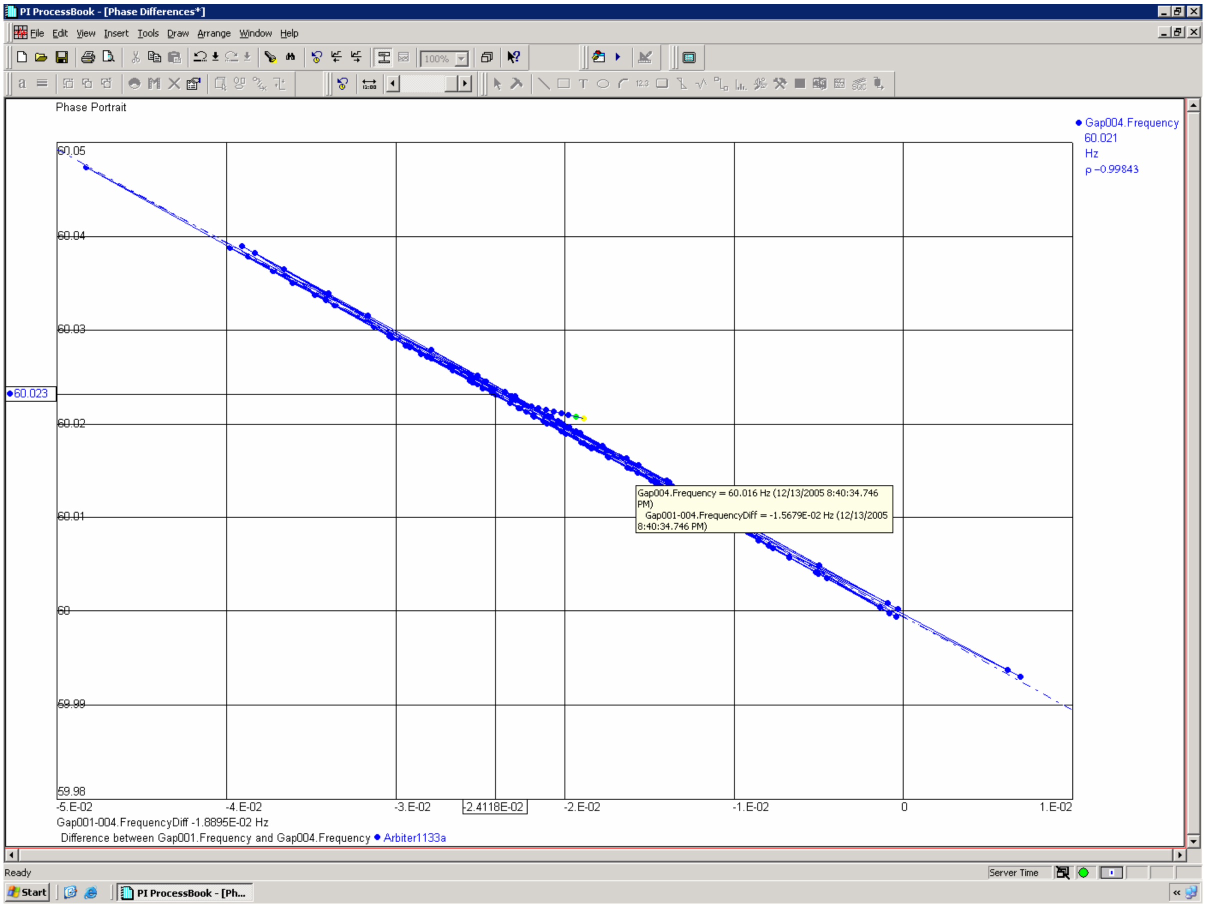

Figure 7 a Phase Portrait of relative angle between Paso Robles and San Leandro

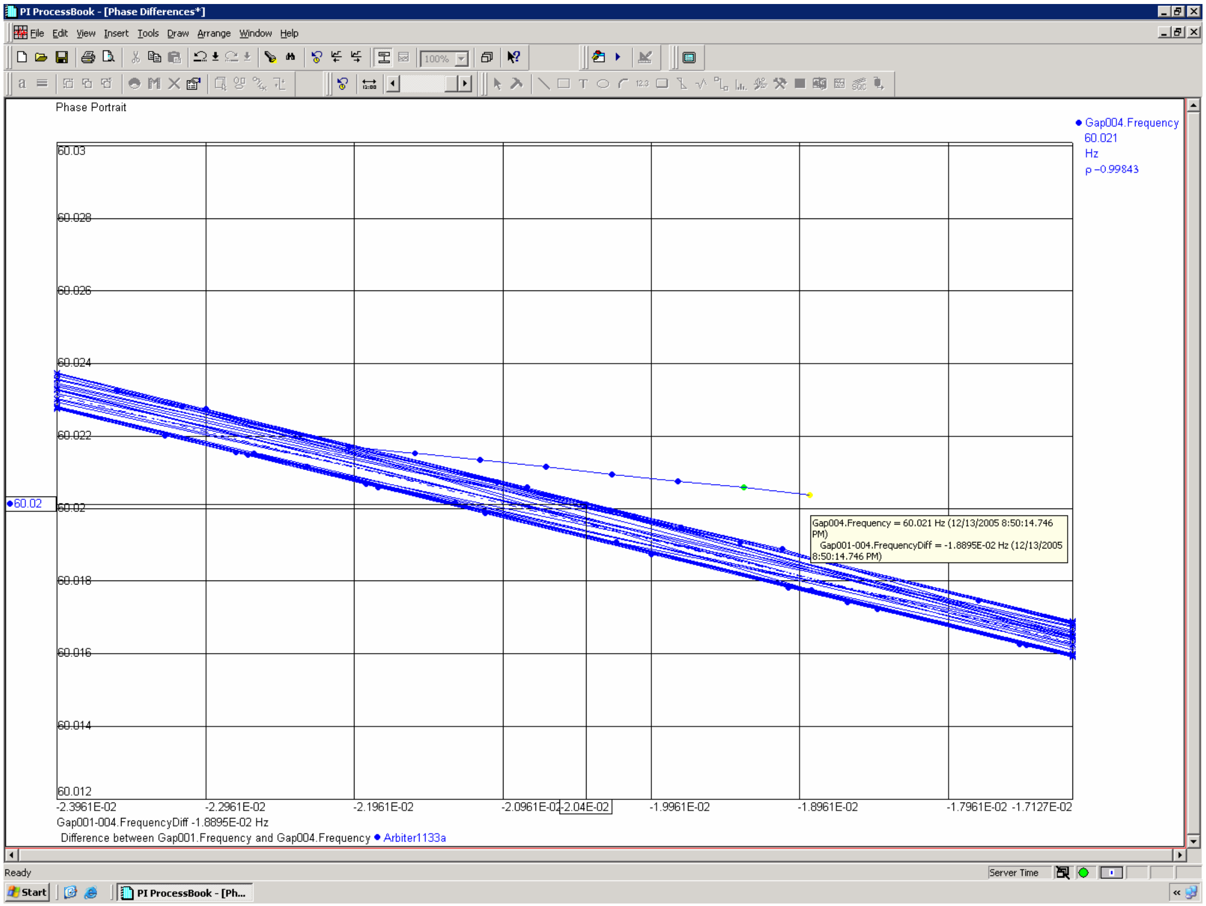

Figure 7 b. Zoomed in view of Phase Portrait.

The cursor is used to mark a specific value of x and y. The mouse-over is used to interrogate specific points including the time of occurrence.

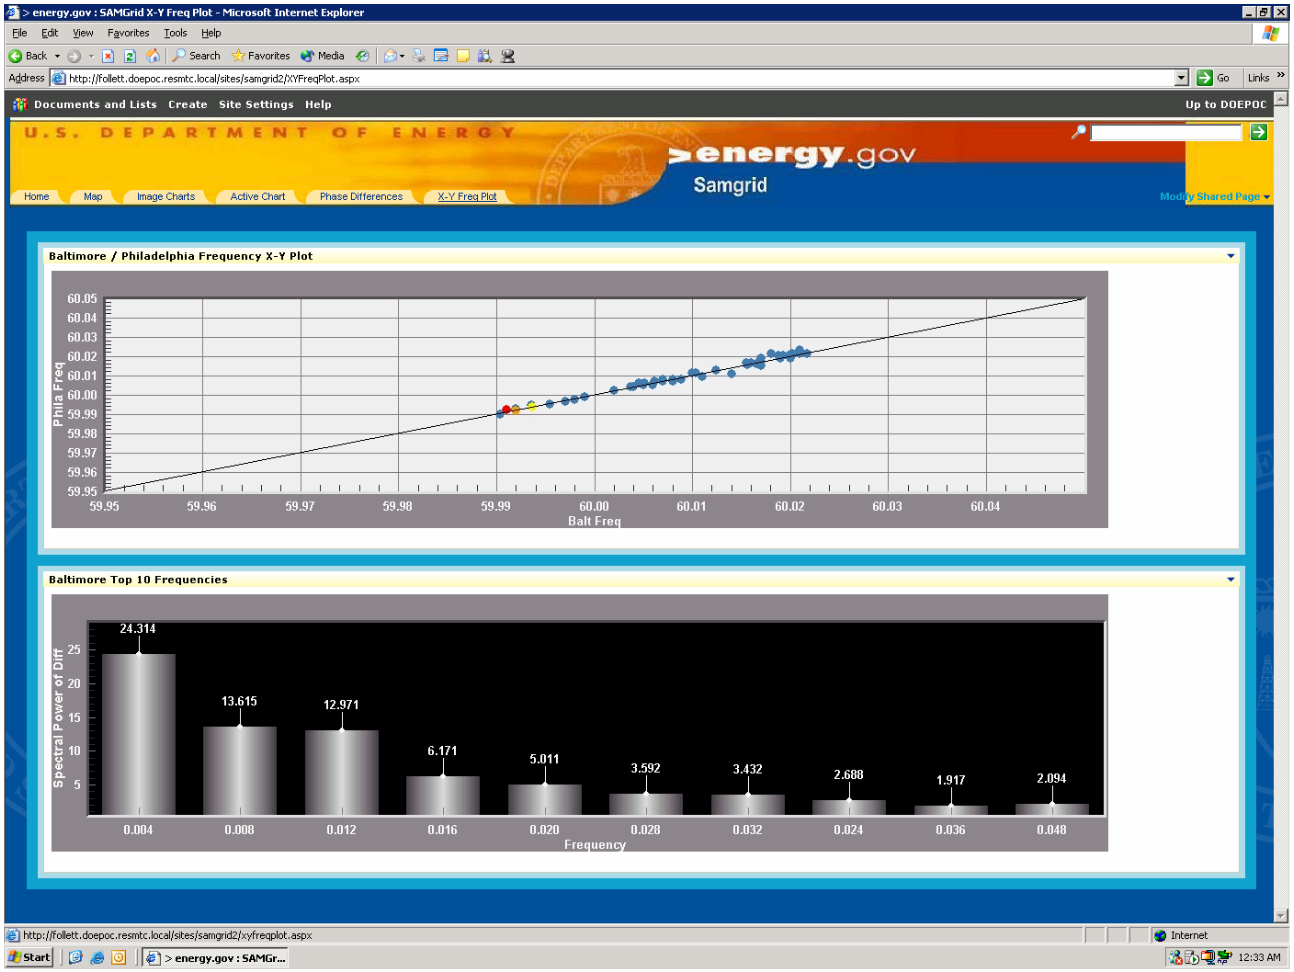

Figure 8. Frequency coherence and FFT peak heights

This is a “thin client” view showing frequency coherence between Baltimore and Philadelphia. These two measurements should be 100 percent correlated in any interconnection. In this case the interconnection is said to be coherent. The bar charts shown in the lower portion of this display represent the peak height of the largest 10 peaks in the Baltimore spectrum.

SQC (Statistical Quality Control)

The SQC chart includes the 8 standard SQC rules and automatically sends alarms when the rules are violated.

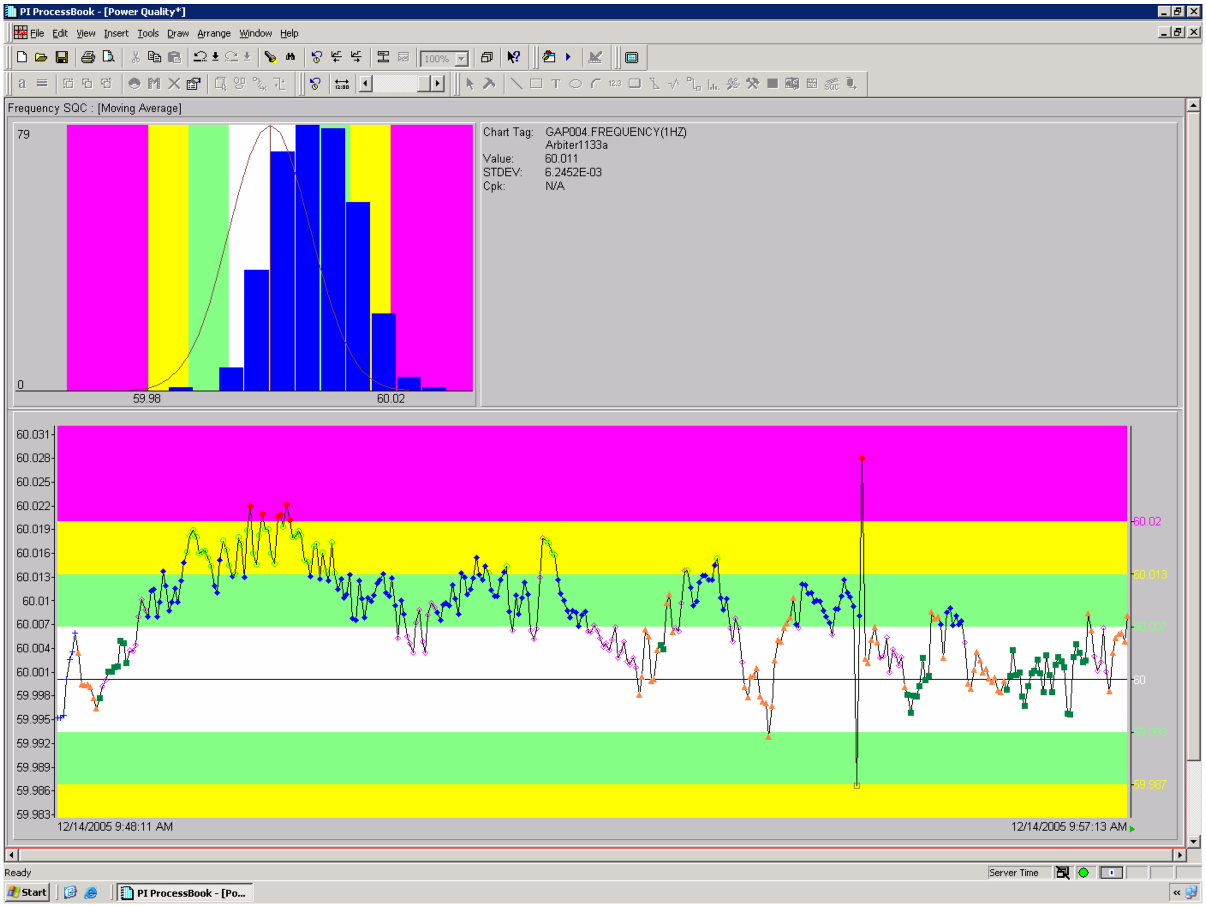

Figure 9 Zoomed in view of SQC chart

This shows the symbol types used to display the SQC rules. Each rule violation is shown using a different shape and colored symbol.

A well known test for a process to be in control is the Akaiki whiteness test. If the process is Gaussian white noise, it is said to be "in control". The well known SQC tests also determine if a process is in control, using a set of ad hoc rules first published in 1929. These rules are included in the PI-SQC package.

We determine the Akaiki test by integrating the area under the FFT curve out to the first 50 harmonic numbers. If the integral is less than 5, the process is said to be "in control".

We run the SQC alarms on the peaks of the FFT of frequency. This clearly indicates when the grid is unstable. These alarms are automatically sent to appropriate staff members via email, LCS, or SMS.

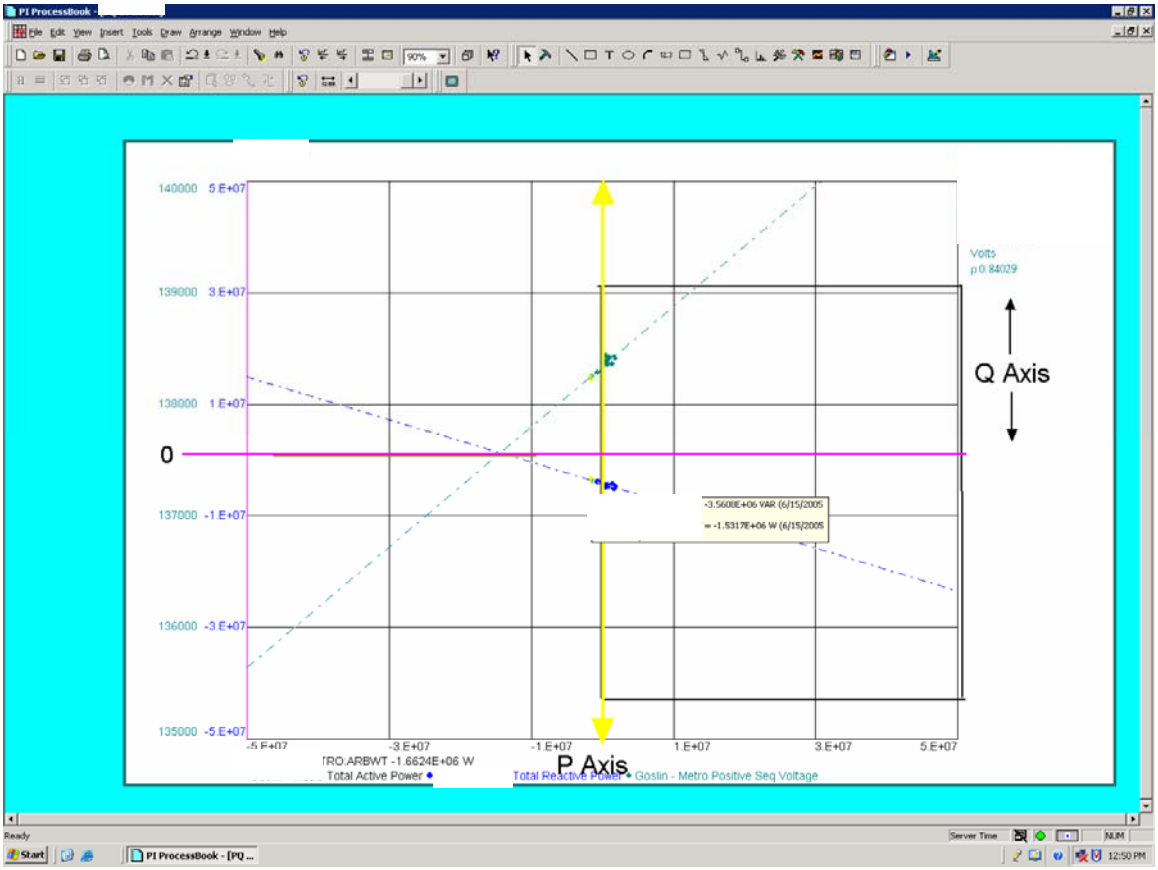

Figure 10 PQ and PV chart

This chart shows real power plotted on the X axis and reactive power and voltage plotted on the Y axis. The current value of voltage and reactive power are indicated by the yellow dot, the next most recent reading is shown by the orange dot. The cursor and mouse-over functions behave as described earlier.

This type of function is used in Wide Area Protection (WAP) relaying. In some relaying systems, the WAP or out of step relays will fire whenever, P-Q exits a unit circle. In this case we observed abnormal behavior when P entered quadrant III.

The PI Grid Monitor provides the means to plot many types of X-Y charts, the P-Q type is one of many that can provide early warning for grid stress. The GM can also provide high level alarms and notifications that will direct an operator's attention to a particular situation.

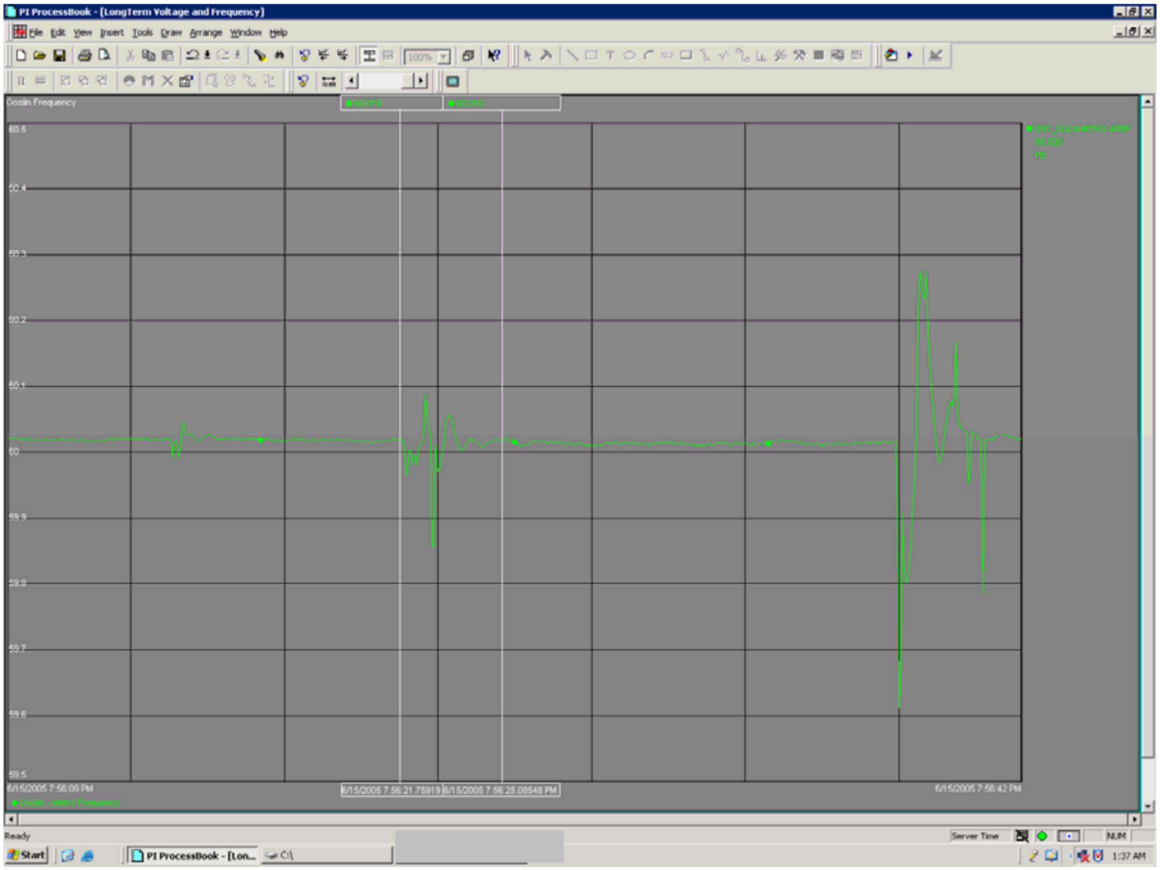

Figure 11 Frequency transients

This shows frequency transients prior to a blackout. Both the amplitude and the duration of the transients are increasing.

Frequency Charts

This is simply a collection of X-Y plots of "A" frequency versus other substations. Users can "double click" on any one and it becomes a full page display with full interactive capabilities.

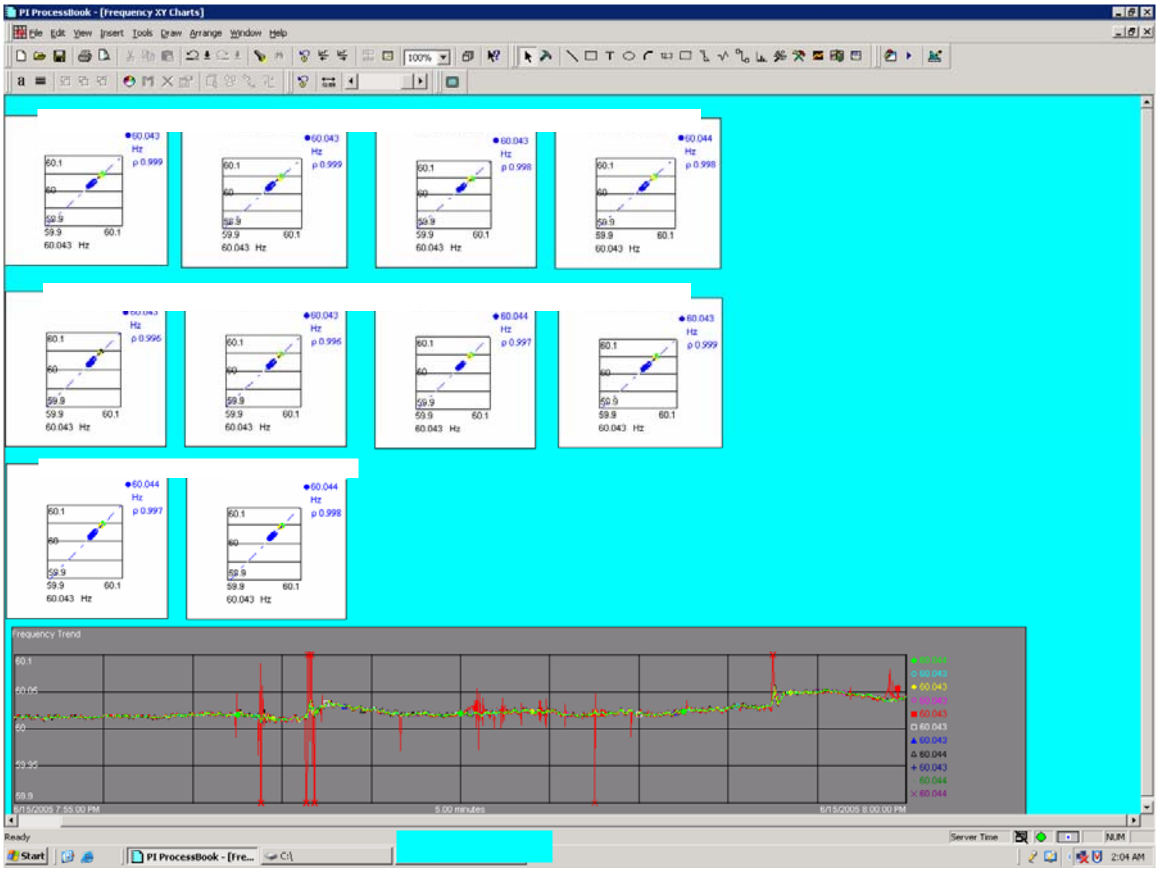

Figure 12. Multiple Frequency coherency charts

A double click on any of the charts causes it to expand to full screen. Each X-Y plot represents a coherency test between substations.

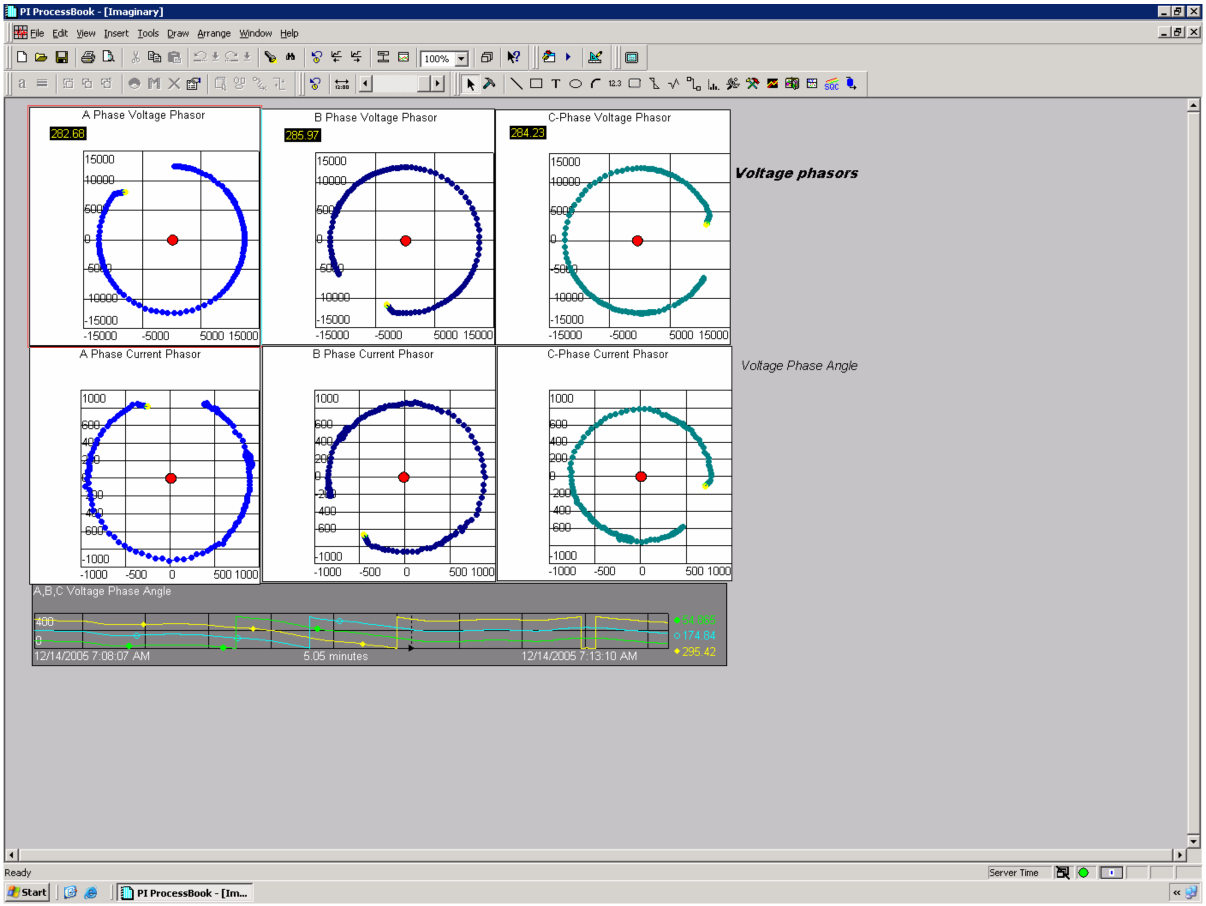

Figure 13. Time histories of phasors

This is a chart showing the time histories of voltage and current phasors. These are X-Y plots of the real and imaginary components of the voltage and current phasors. The current phasor is shown in yellow and the next to the last shown in orange. These are generated at the rate of 20 per second are updated at one second intervals.

The lower curve shows the absolute voltage phase angles.

Conclusions:

We demonstrated that the standard out of the box PI software and ProcessBook are good real time tools for extracting valuable information from PMUs.

The PI system is capable of handling a sustained rate of 80,000 events per second which will support at least 50 PMUs sending data to one server. A typical PMU will output at least 50 measurements at 20 Hz rates, this means that about 1000 events per second will be sent to the archive. Each event contains about 5 bytes on the average, so we are storing about 5K bytes per second, or 432 Mbytes per day without compression. With compression, we commonly reduce the storage by 3/1 so we can assume that each PMU consumes 100 Mbytes per day.

So if 50 PMUs were connected to the PI system, we estimate that 5 Gb per day would be used. So a 2 Tb SAN drive would hold all data from 50 PMUs for over one year.

We demonstrated the used of real time FFTs to determine the real time modes in the grid. This technology can be used to predict blackouts by alarming on the rates of change of the peak heights the FFT spectrum.

We demonstrated the use of SQC charting to provide alerts on grid stability.

We also demonstrated the use of standard PI-ProcessBook X-Y charts demonstrating how to display phase portraits in real time.

We demonstrated how to perform synchronized difference calculations with phasor data and also outlined the discontinuity problem associated with absolute phasor measurement.