Redundancy and Reliability of Wide-Area Measurement Synchrophasor Archivers

Charles H. Wells
OSIsoft, LLC

John Harrell and Roy Moxley
Schweitzer Engineering Laboratories, Inc.

Revised edition released March 2011

Originally presented at the
13th Annual Western Power Delivery Automation Conference, March 2011

Redundancy and Reliability of Wide-Area Measurement Synchrophasor Archivers

Charles H. Wells, *OSIsoft, LLC*
John Harrell and Roy Moxley, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Phasor measurement units (PMUs) are rapidly being installed across the United States as part of the Department of Energy (DOE) Smart Grid Investment Grant (SGIG) Program. Within three years, there are expected to be over 1,000 PMUs installed by nine SGIG grantees. There will be many more PMUs installed in distribution networks to help accommodate intermittent power from rooftop solar panels and wind farms. In addition, PMUs will begin appearing at the terminals of generation equipment, transformers, and large motors.**

Some of the SGIG grantees intend to install highly available redundant measurement systems. One of the key reasons for redundancy is to support North American Electric Reliability Corporation (NERC) requirements to be able to make security patches to software without losing data.

In this paper, we show how fully redundant cybersecure systems can be assembled using standard PMUs and standard commercial software. Using the information available from the fully redundant system, we also show an example of how real-time grid damping calculations can be used to prevent grid collapse.

I. Introduction

Numerous recent events have illustrated the importance and usability of synchrophasor data. Islanding events and system separations have occurred on the North American power grid, where synchrophasors may have been the only system to detect the event. Because of the critical nature of this information, it is not enough to simply transmit data and hope the data arrive at the desired location and are processed correctly. This paper outlines a method of configuring redundant phasor measurement units (PMUs) using commercial off-the-shelf hardware and software. Advanced data analysis is also introduced to increase the usability of the information to operators in order to enhance situational awareness.

II. Recent Events and the State of the Art

In the more than seven years since the formation of the Eastern Interconnection Phasor Project (EIPP) and the subsequent formation of the North American SynchroPhasor Initiative (NASPI), utility companies have begun to understand the value of real-time data from PMUs. This was clearly demonstrated during a storm incident on a major transmission grid [1]. During this event, the role of PMUs in detecting and maintaining a power island was clearly documented, including the published savings of over two million dollars. This is the first documented savings from applying PMUs on the transmission grid. The utility company that documented the savings was one of the first utility companies to widely install commercially available PMUs on both their transmission and distribution grids. The first set of PMUs was installed in 2004 and became fully operational in April 2005. On June 15, 2005, the PMU systems recorded a major blackout. Fig. 1 shows the trend lines of the grid behavior prior to, during, and immediately after a system event.

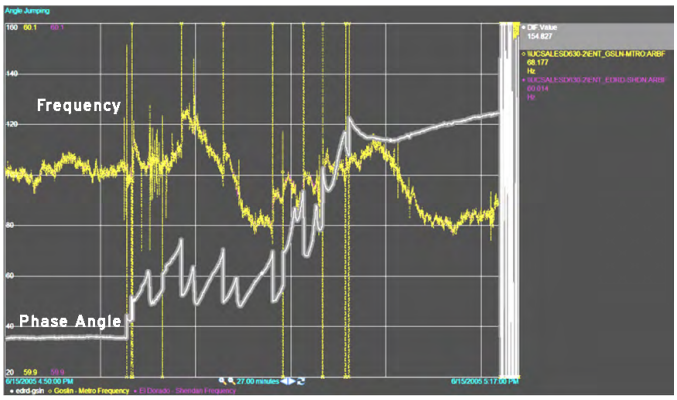

Fig. 1. Trend lines show the behavior of the grid prior to, during, and immediately after a system event

The trend lines in Fig. 1 show the frequencies plotted from two substations more than 400 miles apart (this system is part of the Eastern Interconnection), as well as the angle difference between the two locations. The report shows 27 minutes of data prior to the blackout. The angle difference was very stable all day at around 35 degrees. At about 5:00 p.m., there was a jump in the angle difference. This was subsequently found to be the result of recloser failure on the lines between the two substations. This started the system on an unstable trajectory, which can be seen by the continuing angle increase. The report also shows angle jumps in the negative direction, relieving the stress on the grid temporarily. The angle jumping was due to the fact that, as the lines opened, the impedance between the two substations increased instantaneously because of the line outage (less copper between the two nodes). Similarly, as the lines closed back in, the impedance dropped, resulting in a decrease in the angle difference.

Just prior to the blackout, the angle difference exceeded 120 degrees, clearly an unstable condition [2]. Note that a few minutes later, the grid collapsed. A second analytical view of the data is shown in Fig. 2. This is a continuous plot of the grid frequency at each bus versus the frequency at the reference bus.

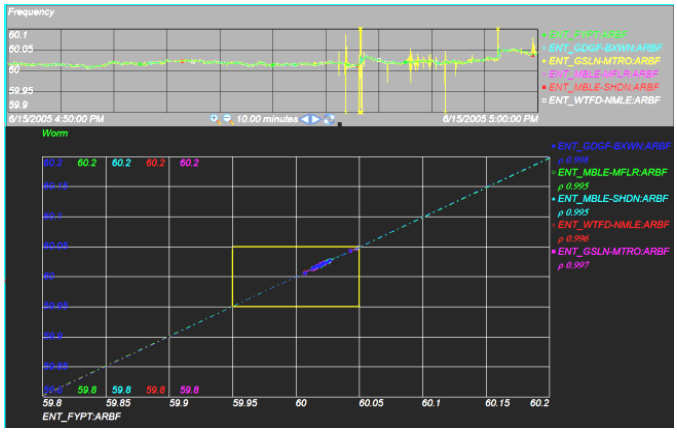

Fig. 2. Overlays of frequency and frequency x-y plots

In Fig. 2, the x axis shows the frequency at one of the substations, and the y axis shows the frequency at the other buses in the network. The dots on the chart move up and down the 45-degree line, showing that the grid is coherent. The labels of the variables and the correlation coefficient of the line are shown on the right-hand side.

Another analytical view that must be computed is the angle difference between parts of the grid, as shown in Fig. 1. We show a more recent example of relative angles in the Eastern Interconnection in Fig. 3.

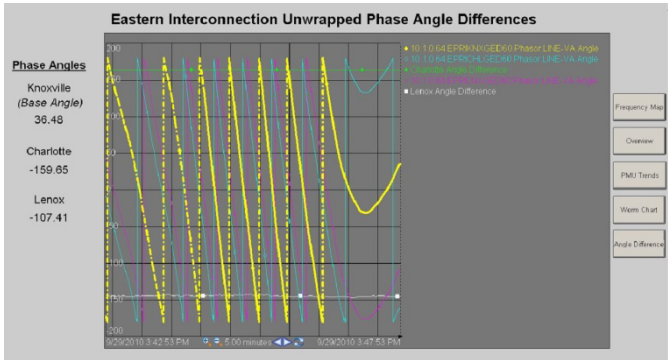

Fig. 3. Unwrapped angle differences – multiple stations

Fig. 3 shows three absolute phase angles and two relative phase angles. The A-phase in each of these areas is different, which is the reason for the large angle separations. The absolute phase angles are different by approximately 120 degrees; however, it is the relative angle that is the stability indicator. The absolute angles wrap from –180 degrees to +180 degrees. The frequency is lower than 60 Hz for the first part of the trend and is exactly at 60 Hz when the absolute angle slope is zero. When the rate of change is positive, the frequency is above 60 Hz. This can be seen on the right-hand side of the trend chart. The relative angles between Lenox-Knoxville and between Charlotte-Knoxville are also shown on the trend chart. The relative angles are called the unwrapped angle difference. Fig. 4 shows an example of how the C++ code for the unwrapping could be done using a very simple algorithm.

```
static double unwrapped_angle_difference
(double A, double B)
{
      double Diff = A - B;
      return
          Diff < -180 ? Diff + 180
        : Diff > 180 ? Diff - 180
        : Diff
        ;
}
```

Fig. 4. C++ code for unwrapping angle differences

This calculation can only be achieved with time-aligned data. The alignment can be completed in phasor data concentrators (PDCs) or directly in the software server. In this case, there are only three data streams that have to be aligned to perform the unwrapped angle calculation.

As the importance of phasor measurements (demonstrated by Fig. 1, Fig. 2, and Fig. 3) becomes better known to utility companies, there will be a rapidly increasing dependence on these measurements. No longer are PMUs considered laboratory instruments. Synchrophasors are part of critical control systems. Other utilities are using phasor measurements in real-time control as well. Hence, the industry needs a secure and highly available system to provide these measurements. We outline how this can be accomplished using standard off-the-shelf software and hardware to create a cybersecure synchrophasor platform (CSSP).

III. Reliability and Availability

In this section, we discuss how to increase the availability of synchrophasor measurements by creating a CSSP using redundant, reliable systems.

A. Reliability Calculations

A number of reports and papers have been written on redundancy for protective relay systems, such as [3], [4], and [5]. For example, the IEEE report "Redundancy Considerations for Protective Relaying Systems" "provides the relay engineer with information about what factors to consider when determining redundancy requirements … [and] addresses differences depending on application area, present practices and provides real world examples" [3].

Reference [4] uses fault tree analysis to show that, for permissive overreaching transfer trip (POTT) schemes, dual redundancy has higher dependability and better security and the incremental cost is a low price to pay (where security is defined in the classical protection manner to be probability of an incorrect operation). Unavailability calculations from [6] show that a dual-redundant common-mode system unavailability is ten times lower than the unavailability of a single relay system. Reference [6] concludes that dual-primary protection from the best manufacturer is the best design choice.

In this paper, we discuss reliability and redundancy for synchrophasor archiver systems. All computer systems consist of the application, operating system (OS), and hardware.

In the CSSP outlined in this paper, the application is a high-availability historian software server, the OS is Windows Server® 2008, and the hardware is a substation-hardened computer.

According to the "ITIC 2009 Global Server Hardware and Server OS Reliability Survey," Windows Server 2008 running on typical Intel®-based platforms has an unplanned annual downtime of 2.42 hours [7]. Windows Server 2003 was slightly higher, with an annual downtime of 3.02 hours [7].

Given a combined OS and typical hardware downtime of 2.42 hours, we can easily determine that availability is 99.97 percent.

The selected hardware has a field-proven mean time between failures (MTBF) of much greater than 100 years. This means that for every 100 units installed, there is one or fewer failures per year. Assuming a worst-case scenario (where the unit has to be shipped to the manufacturer for repair), the availability of the substation-hardened computer becomes 99.99 percent. By using the substation-hardened computer, the combined OS and hardware availability number is higher because the hardware platform has over ten times the reliability of typical industrial computers and a much higher reliability than typical office-grade computers.

To achieve higher availability, the CSSP is placed into a highly available architecture. When one system fails, the other system is available and continues operation.

The historian software uses a concept known as a collective (a common name for a collection of servers). All computers in the collective have identical data with connections managed by the client software. For example, when a new client connection is requested, the connection is made to the least-loaded computer, providing the fastest response.

The availability of two systems in parallel is determined by the following equation:

$$1 - ([1 - \text{availability}] \bullet [1 - \text{availability}]) \quad (1)$$

With an availability of 99.97 percent for the combined Windows Server 2008 and a typical Intel-based platform, we obtain a mirrored system availability of 99.99999 percent for the CSSP.

B. Data Loss and System Maintenance

Patch management has become such an important part of managing servers that research dedicated to just this topic is being performed [8] [9]. These studies show that there is a significant amount of downtime necessary to patch a server OS. Without redundancy, the downtime necessary to patch a server OS is the time span for loss of synchrophasor data. Both studies included the Windows Server 2003 OS. Windows Server 2008 was not included in these studies presumably because there are not enough field data at this time. Windows Server 2008 includes specific improvements in the OS that affect reliability, availability, and serviceability. Specifically, hot patching has been introduced to reduce the number of times a reboot is necessary to finalize a patch installation.

North American Electric Reliability Corporation (NERC) CIP-007-2a Cyber Security – Systems Security Management (R3, Security Patch Management) states that "the Responsible Entity…shall establish, document and implement a security patch management program" [10]. As described by NERC in [11], there is a need to architect and design systems that have a commensurate level of availability. NERC states specifically that implementation should be done securely in redundant pairs to avoid systemic data gaps while standard maintenance is performed on the system.

This is why we call the solutions presented here "cybersecure." The NERC patches can be made on each machine at any time with no loss of data.

IV. Cybersecure Synchrophasor Platform

In this section, we describe the hardware component of the system.

A. PMUs

The CSSP requires two PMUs that can each send identical User Datagram Protocol (UDP) packets to the software. The header and command packets can be sent via Transmission Control Protocol/Internet Protocol (TCP/IP) because these packets are sent only at limited rates. However, the data streams flow at high rates, often including more than simply phasors. For example, in one wide-area measurement system, 22 PMUs are used (soon to be expanded under the Smart Grid Investment Grant Program). Each PMU sends 84 measurements in each packet. The data include phasor vectors in a rectangular or polar format, frequency, rate of change of frequency, and other measurements, including positive-sequence voltage and current phasors.

Header commands are sent by the interface to the PMUs every minute to determine if the PMUs have been reconfigured within the last minute. This information is part of the IEEE C37.118 standard, and any compliant PMU includes it.

B. Computers

The CSSP uses substation-hardened computers. The hardware is based on low-power mobile-class central processing units (CPUs) configured to run the Windows Server 2008 standard 64-bit operating system. The computers are rated for high- and low-temperature operations, require no fans, and use 2 GB of memory with 120 GB of solid-state drive (SSD). There are no moving components in this system.

C. IEEE C37.118 Software Interface

A software application that accepts IEEE C37.118 messages and converts them to OPC format, referred to as an interface, is loaded on each substation-hardened computer. This conversion allows the data collected through the IEEE C37.118 messages to be available to other software applications on the same computer that understand OPC but not IEEE C37.118. There are two instances on each machine

running simultaneously, one for PMU A and the other for PMU B. Each of these machines is configured to collect multiple data streams from each PMU. In addition to the data collected via the IEEE C37.118 messages, 600 measurements are collected once per second from individual PMUs. These measurements include harmonic values in each phase for both current and voltage. This is especially important when using the CSSP for transformer condition monitoring.

For dual-redundant failover, there is a dual interface instance for every PMU, with a heartbeat to manage the failover in each interface. The heartbeat runs at 30 Hz. While there can be latency during failover, no data loss occurs because each dual interface instance exactly mirrors or buffers the PMU data.

The system is configured to failover to a secondary interface when the primary interface fails. The interfaces on both computers run in parallel. One is considered Primary A, and the other is Primary B because they both have full and equal capabilities, yet only one is active at a time. This is often referred to as dual primary. When a failover occurs, the alternate system becomes the active primary and the original interface is inactivated. When the original interface is again available, it automatically becomes the alternate primary, so it can become the active primary again when the alternate fails.

D. Software System

The high-availability historian software server is basically two identical historian software servers running on separate substation-hardened computers. In this case, one historian software server runs on Computer A, and another historian software server runs on Computer B.

The two computers form a collective. Clients, when connecting to the server, use the collective name. They do not explicitly know to which historian software server they are connected. The client software requests connections to the collective manager. Clients can force a connection to a specific server; however, if that server fails, the connection is passed to the alternate server. The software architecture is shown in Fig. 5.

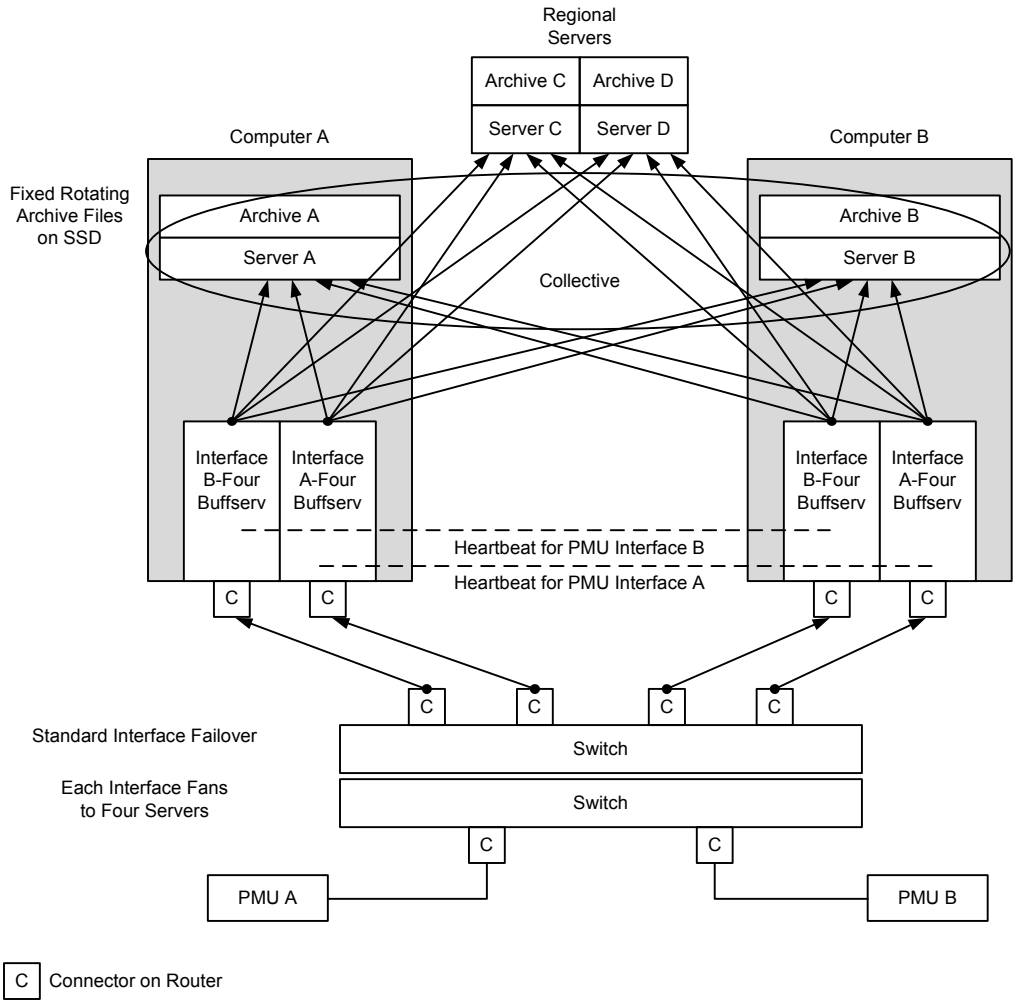

Fig. 5. Sketch of the software architecture

E. Display

The display of phasor data requires either a desktop client or a web-based application. Display software from most manufacturers includes tools to simplify selections of PMU and data type to display.

Most display systems include a method to jump or index from one PMU to another using a menu or drop-down box. One display example is shown in Fig. 6.

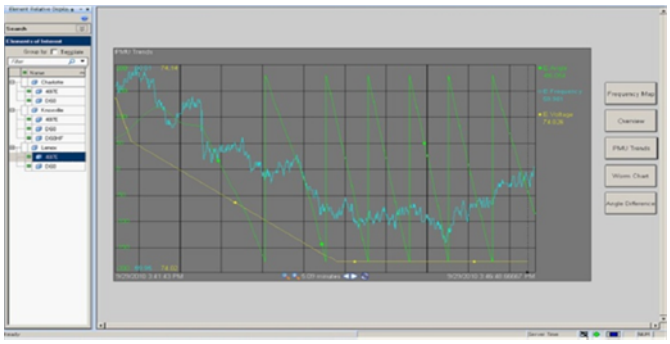

Fig. 6. Element relative display

Fig. 6 shows one of the seven PMUs that were operating on September 29, 2010. The user can select any PMU device listed on the left to see the detailed trend lines for that device. The trend can be panned and zoomed to show history and details.

V. Advanced Data Analysis

A. Angle Difference Comparisons

There are large differences when comparing the results of computing angle differences on a synchronous versus asynchronous basis. Fig. 7 shows an example of the differences.

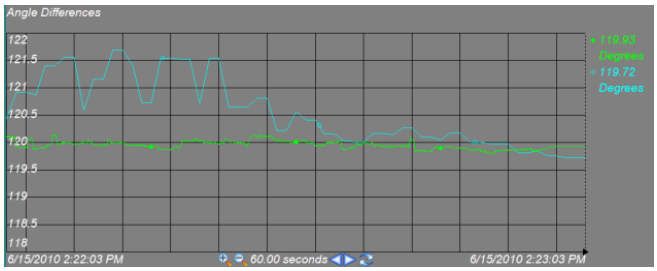

Fig. 7. Angle difference comparison

Fig. 7 shows the asynchronous values of the angle difference in the top line. The synchronous calculations are shown in the bottom line. The time alignment is performed on the computer at a 30 Hz rate. The estimated latency in this example is less than 3 milliseconds.

B. Other Real-Time Analytics

Early control system analysis was done using second-order systems [12]. Modern power systems are too complex to be completely constructed using second-order systems. It is useful to construct a model made up of a number of second-order components, as has been done, starting with generators [13]. We can perform a second-order analysis at a number of frequency points to gain insight on system stability. Consider a classical second-order system with *y* as the input and *x* as the output. The *x* variable is the time-synchronized frequency difference between two disparate buses in the region under analysis. The *x* variable is computed at a 30 Hz rate from PMUs located at the two buses. In this case, a second-order system is selected for computational simplicity and provides useful visualization and damping calculations. A second-order system describes an RLC circuit, such as can describe transmission lines. A higher-order characterization would provide a more exact model, but only marginally more useful information would be extracted.

It is likely that in the near future there will be PMUs at every major bus in the network; hence, it is possible to compute all combinations of frequency differences across the region.

From classical dynamics of second-order systems, the system model can be written as follows:

$$K(s) = K_0 \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \tag{2}$$

where:

$s = j\omega$

$$K(\omega) = \frac{1}{1 - \left(\frac{\omega}{\omega_n}\right)^2 + j2\zeta\left(\frac{\omega}{\omega_n}\right)} \tag{3}$$

A plot of the transfer function using frequency domain analysis results in the shape shown in Fig. 8 for any particular damping factor and frequency.

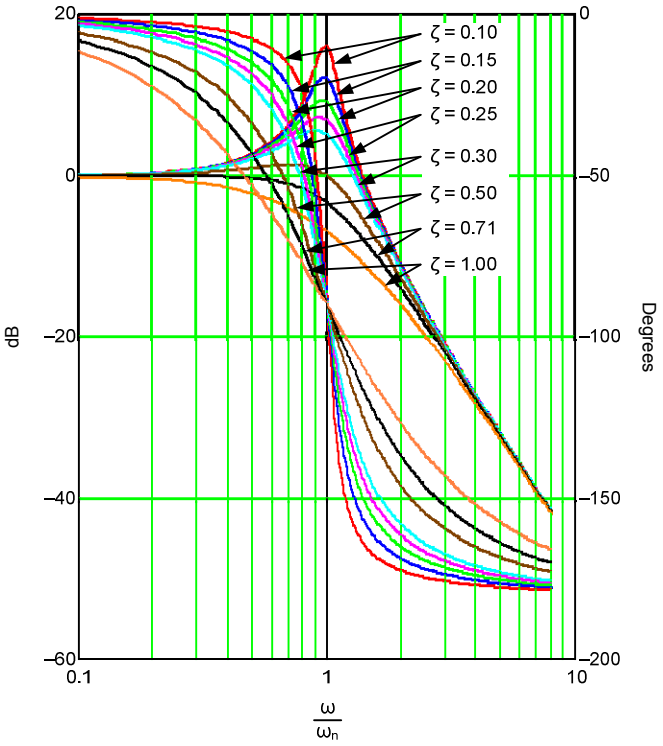

Fig. 8. Frequency analysis plot of a second-order system

For systems that are critically damped, the damping factor (ζ) is equal to 1.0. For underdamped systems, the damping factor is less than 1.0, and a resonant peak occurs at the natural system frequency. As the damping factor approaches 0, the system is about to become unstable [14]. The poles of the system are located on the imaginary axis in the *s* plane.

Input variables to the actual power system network are large, random perturbations that contain a large number of sinusoidal inputs. Thus the output contains a large number of sinusoidal components that can be extracted using Fourier analysis. We can identify the transfer function in real time using the Fourier transform method.

In a power network, we expect to see one or more resonant frequencies. These are caused by the network itself, as well as poorly tuned generation equipment "hunting" against each other. A classical case of this is shown in [15], as well as [16]. However, these papers describe the properties of the oscillation after it has started rather than predicting that the system could go unstable.

The following sections demonstrate how this method works on actual data prior to a major separation event.

1) Raw Data

A major grid separation event occurred, and substation PMUs collected synchrophasor data prior to, during, and after the event. Frequency trends from two stations are shown in Fig. 9. Station A is outside of the power island event, and Station B is inside the power island. The frequencies track each other until just prior to the event. This display is what an energy management system operator might see. The solid line is outside the power island, and the dashed line represents the frequency inside the power island.

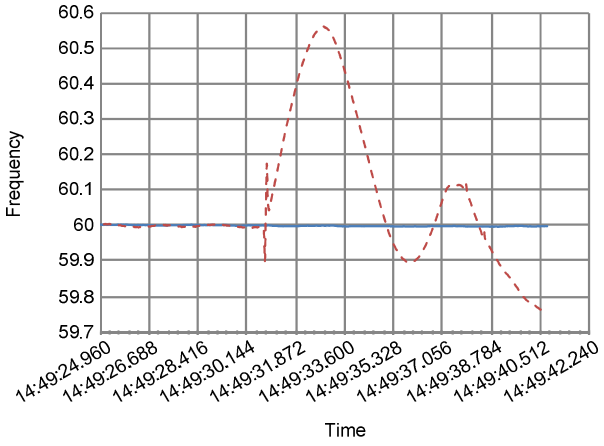

Fig. 9. Separation event

The synchronized differences between the two stations result in the data shown in Fig. 10.

There was very little indication of a problem with the grid from the raw frequency differences. There are two small disturbances at about 2 minutes and 1 minute prior to separation. But these types of disturbances are common and often occur as a result of a line reclosing or several attempts at a reclose.

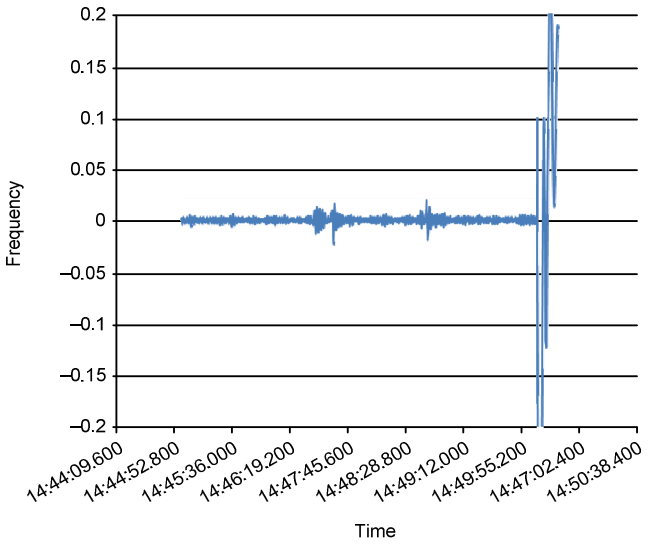

Fig. 10. Frequency difference between stations

2) Plots of Frequency Inside and Outside Before Separation

We used the Fast Fourier Transform (FFT) to compute the frequency domain plot for each station. Station A is outside of the eventual power island and about 200 miles from the center of the island. A plot shows the relationship amplitude of the FFT versus spectral frequency plotted on log-log paper, as shown in Fig. 11.

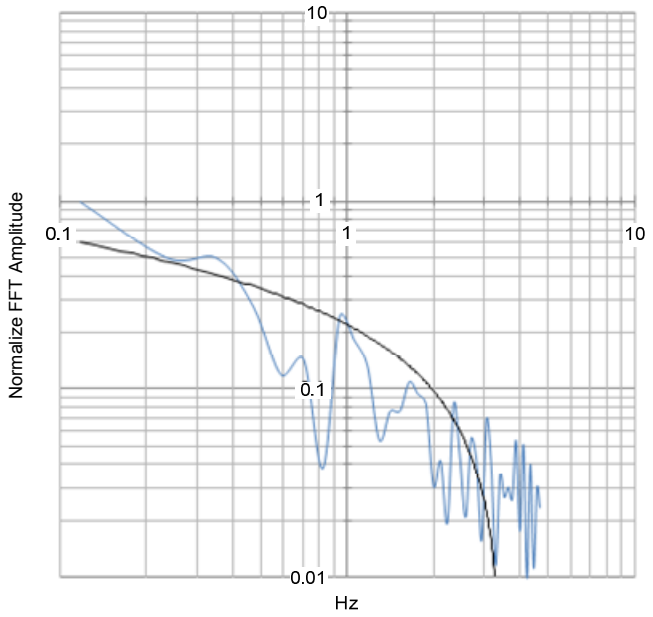

Fig. 11. Frequency domain plot of Station A frequency response

The system behaves as if it were overdamped, and there are no significant oscillations around Station A (i.e., a system with a damping ratio greater than 1.0, commonly known as a stiff system). The region around Station A was normal 4 minutes before the separation occurred. The thin solid line is a logarithm trend line fitted to the spectral data.

A plot of the Station B response is shown in Fig. 12.

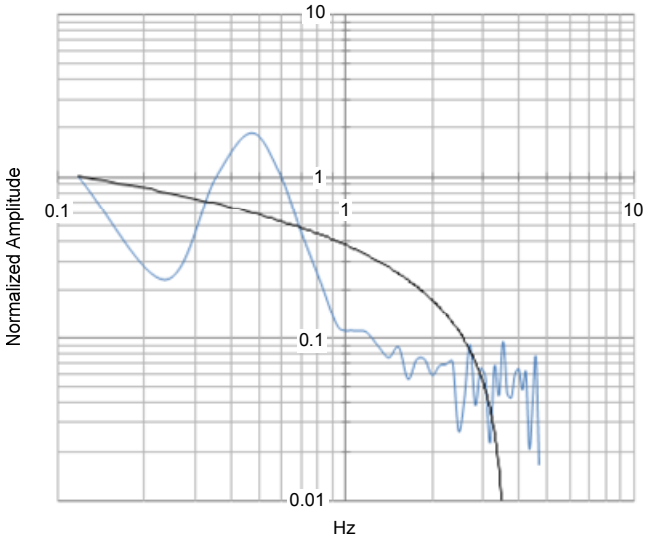

Fig. 12. Station B frequency response

Note there is an indication that the damping at 0.48 Hz inside the power island is less than 1.0. Also, the maximum amplitude is less than 2.0. This method of grid instability detection is described in [17].

3) Plot of the Frequency Differences

A frequency domain plot of the frequency response across the region is shown in Fig. 13.

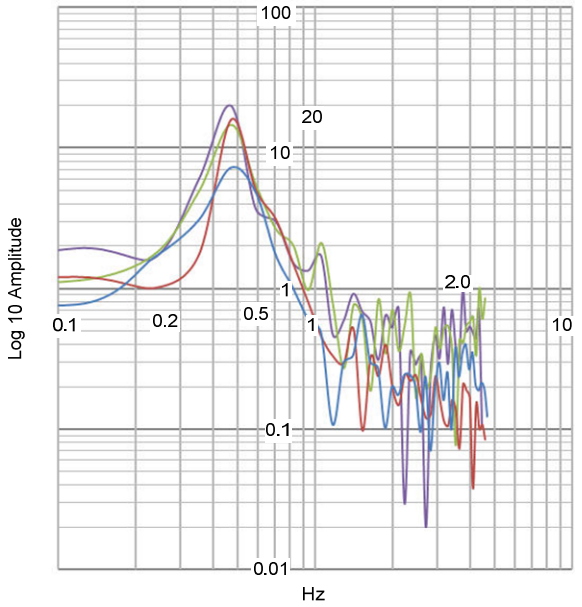

Fig. 13. Plot of frequency difference at 4, 3, 2, and 0 minutes before separation

The resonant peak is located at 0.468 Hz and continually grows from about 7 to 20 in the 4-minute interval before the separation. Comparing this with the frequency response of Station B, it is clear that the peak amplitudes are significantly larger and, had these been available in real time, could have provided an early warning to the operators of an impending separation.

The question arises as to what the frequency response looks like during normal, stable periods. Fortunately, there were data collected in June 2005 for the frequency measurements at these two stations. The plot is shown in Fig. 14.

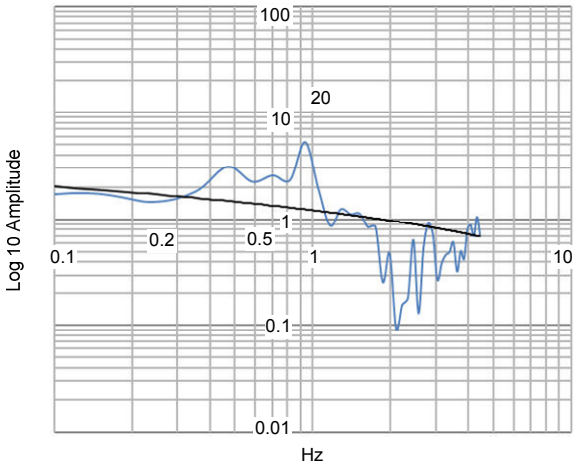

Fig. 14. Frequency response between stations

In this case, the grid is much more stable, with a small resonant peak at 0.95 Hz and a very small peak at 0.48 Hz. This clearly shows a small signal stability problem that has existed for more than four years, although it was not nearly as pronounced as during the major grid separation event.

4) Damping Coefficient

One nontraditional way to represent the damping coefficient may be computed directly from the FFT amplitude at the resonant peak. In this method, the damping term (ζ) may be computed directly from (4). The resulting formula is:

$$\zeta = \frac{A_0}{2A_i} \tag{4}$$

where:

A_0 is the amplitude of the FFT at harmonic number 0 (for normalized FFTs, this value is always 1).

A_i is the amplitude of the FFT at harmonic number *i*.

The value of *i* is the location of the resonant peak.

This method uses a sliding sample window to compute a time domain quantity (damping) from a frequency domain value (FFT amplitude). When an excited mode includes a decaying dc offset, the ratio of the mode energy (A_k) to the dc energy (A_0) approximates the damping of the mode. Note that the formula for converting from harmonic number to hertz is as follows:

$$\text{Hz} = \frac{i}{N\Delta t} \tag{5}$$

where:

N is the number of samples in the moving window.

Δt is the sampling interval.

In the case shown in Fig. 7, the value of *i* at the resonant peak is 5, with a sampling interval of 0.033333 seconds, and N is equal to 256.

A plot of the damping coefficients at 0.468 Hz is shown in Fig. 15.

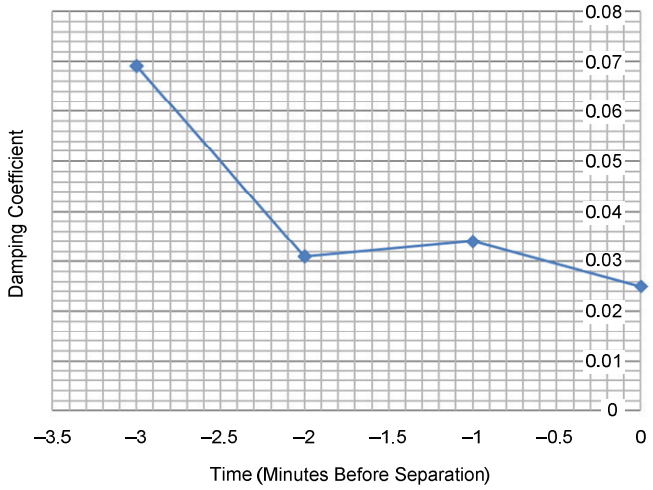

Fig. 15. Damping coefficients prior to separation

The damping is very low at the beginning of the plot and continues to drop to 0.025 just prior to the separation. This implies that a moving window FFT with 256 points in the window at 0.033333 time intervals between the samples is a reasonable starting point for the grid failure detection method. This is about an 8.5-second moving window.

VI. Conclusions

Recent events have shown that synchrophasor data are useful to operators of electric power systems. This paper details a fully redundant cybersecure PMU system assembled using commercially available off-the-shelf products. Using the information available from these redundant systems, new analytics can be used to increase the reliability of the power grid. The topics discussed in this paper demonstrate or illustrate the following points:

- Wide-area management system projects funded by the U.S. Department of Energy are encouraged to use commercially available products. These products are available and applied today.
- Regulatory groups are defining redundancy requirements.
- Synchrophasor data can be displayed using direct measurements or calculated values to improve operator understanding of events.
- As applications expand, a flexible system to view calculation results is desirable.
- Available information can be used to increase the reliability of the power grid with new analytics.

Just as relay systems are designed to be secure and reliable, synchrophasor systems apply redundancy techniques to achieve these goals.

VII. References

[1] F. Galvan, S. Mandal, and M. Thomas, "The Role of Phasor Data in Emergency Operations," *Transmission & Distribution World*, December 2008. Available: http://tdworld.com/overhead_transmission/role_phasor_data_emergency_operations_1208/.

[2] M. Ilić and J. Zaborszky, *Dynamics and Control of Large Electric Power Systems*. John Wiley & Sons, Inc., 2000.

[3] IEEE Power System Relaying Committee, Working Group I-19, "Redundancy Considerations for Protective Relaying Systems," 2010. Available: http://www.pes-psrc.org/.

[4] E. O. Schweitzer, III, D. Whitehead, H. J. Altuve Ferrer, D. A. Tziouvaras, D. A. Costello, and D. Sánchez Escobedo, "Line Protection: Redundancy, Reliability, and Affordability," proceedings of the 37th Annual Western Protective Relay Conference, Spokane, WA, October 2010.

[5] NERC System Protection and Control Task Force, "Protection System Reliability: Redundancy of Protection System Elements," November 2008. Available: http://www.nerc.com/docs/pc/spctf/Redundancy_Tech_Ref_1-14-09.pdf.

[6] H. J. Altuve Ferrer and E. O. Schweitzer, III (eds.), *Modern Solutions for Protection, Control, and Monitoring of Electric Power Systems*. Schweitzer Engineering Laboratories, Inc., Pullman, WA, 2010.

[7] Information Technology Intelligence Corp., "ITIC 2009 Global Server Hardware and Server OS Reliability Survey," July 2009. Available: ftp://public.dhe.ibm.com/common/ssi/ecm/en/pol03058usen/POL03058USEN.PDF.

[8] Wipro Technologies Product Strategy & Architecture Practice (PSA) Reliability Benchmarking Series, "A Lab-Based Comparison Study of Windows and Linux Enterprise Servers, Study #1: Server OS Patch and Update," March 2008. Available: http://www.microsoft.com.

[9] L. DiDio, "Unix, Linux Uptime and Reliability Increase; Patch Management Woes Plague Windows," Yankee Group Research, Inc., January 2008. Available: http://www-07.ibm.com/systems/kr/power/software/pdf/polo3003usen.pdf.

[10] NERC CIP-007-2a Cyber Security – Systems Security Management, 2009. Available: http://www.nerc.com/files/CIP-007-2a.pdf.

[11] NERC, "Real-Time Application of Synchrophasors for Improving Reliability," October 2010. Available: http://www.nerc.com/docs/oc/rapirtf/RAPIR%20final%20101710.pdf.

[12] J. C. Maxwell, "On Governors," Proceedings of the Royal Society, No. 100, 1868.

[13] Z. Xu, W. Shao, and C. Zhou, "Power System Small Signal Stability Analysis Based on Test Signal," proceedings of the 14th Power Systems Computation Conference (PSCC), Sevilla, Spain, June 2002.

[14] C. H. Wells, "Power Grid Failure Detection System and Method," U.S. Patent 7490013, February 2009.

[15] G. Kobet, R. Carroll, R. Zuo, and M. Venkatasubramanian, "Oscillation Monitoring System at TVA," North American SynchroPhasor Initiative (NASPI) Work Group Meeting, Chattanooga, TN, October 2009. Available: http://www.naspi.org/meetings/workgroup/2009_october/presentations/kobet_tva_oscillation_monitoring_tools_20091008.pdf.

[16] S. Kolluri, S. Mandal, F. Galvan, and M. Thomas, "Island Formation in Entergy Power Grid During Hurricane Gustav," proceedings of the IEEE Power and Energy Society General Meeting, Calgary, Canada, July 2009.

[17] C. H. Wells, "Non-Linear Observers in Electric Power Networks," U.S. Patent 7498821, March 2009.

VIII. BIOGRAPHIES

Charles H. Wells was born in Kokomo, Indiana. He graduated from Vanderbilt University with a BS in chemical engineering and from Washington University with an MS in chemical engineering and a PhD in electrical engineering. Mr. Wells is a registered professional engineer in both chemical engineering and control engineering in the state of California. His employment experience includes Systems Control, Measurex, EPRI, and OSIsoft. His special fields of interest include control systems. Mr. Wells has published over 60 technical papers and is a coauthor of two textbooks. He has been granted nine patents.

John Harrell received his BS in mathematics at California Polytechnic State University in 1984 and his MS in computer engineering at Santa Clara University in 1991. He has worked in a variety of engineering and marketing management and executive positions for companies including IBM, Sun Microsystems, and Sybase. Mr. Harrell's engineering work includes large parallel digital signal processor (DSP) firmware development and distributed control systems. He started working at Schweitzer Engineering Laboratories, Inc. in 2007 as product manager for computing systems.

Roy Moxley has a BS in electrical engineering from the University of Colorado. He joined Schweitzer Engineering Laboratories, Inc. (SEL) in 2000 and serves as marketing manager for protection products. Prior to joining SEL, Mr. Moxley was with General Electric Company as a relay application engineer, transmission and distribution (T&D) field application engineer, and T&D account manager. He is a registered professional engineer in the state of Pennsylvania and has authored numerous technical papers presented at national and international relay and automation conferences. Mr. Moxley also has a patent for using time error differential measurement to determine system conditions.