# Case Study: High-Availability Secure Microgrid Monitoring System for UCSD


Charles H. Wells and Gary Siefert
*OSIsoft, LLC*

William Torre and Byron Washom
*University of California, San Diego*

Ian Olson
*Schweitzer Engineering Laboratories, Inc.*

John Harrell
*formerly of Schweitzer Engineering Laboratories, Inc.*



Presented at the
Power and Energy Automation Conference
Spokane, Washington
March 26–28, 2013

# Case Study: High-Availability Secure Microgrid Monitoring System for UCSD

Charles H. Wells and Gary Siefert, *OSIsoft, LLC*
William Torre and Byron Washom, *University of California, San Diego*
Ian Olson, *Schweitzer Engineering Laboratories, Inc.*
John Harrell, *formerly of Schweitzer Engineering Laboratories, Inc.*

***Abstract*—This paper describes a cybersecure synchrophasor platform (CSSP) that uses standard rack-mounted hardware and software in a single bay rack. It is installed at the University of California, San Diego (UCSD).**

**The CSSP consists of redundant phasor measurement units (PMUs) with dual substation-hardened computers in the UCSD microgrid system. The purpose of the CSSP is to provide an advanced warning of stress developing in the area power system, allowing UCSD operators to prepare for a potential island maneuver.**

**The CSSP includes archives of PMU synchrophasor data and real-time computations of Fast Fourier Transforms (FFTs) and damping coefficients. Moving window FFTs are computed at a rate of 30 Hz, and the damping coefficients at each mode are displayed and archived. The angle separations and FFTs of the data are used to identify a potential island maneuver that might be executed.**

**The data from the CSSP are used for both real-time display and the main campus data archiving system. Data are collected directly from the wall outlet, and voltage, frequency, and phase angles are displayed in real time.**

**With dual substation-hardened computers, the CSSP has redundant failover capabilities, yielding high availability. This configuration allows the system to be patched per the North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) standards while still continuing to collect data, thereby satisfying NERC CIP requirements.**

**The principal results of this paper include data demonstrating the value of high-speed synchronized data collection systems in microgrids.**

## I. Introduction

A cybersecure synchrophasor platform (CSSP) is installed at the University of California, San Diego (UCSD). The CSSP collects data from phasor measurement units (PMUs) and includes archives of data and real-time computations of Fast Fourier Transforms (FFTs) and damping coefficients. This paper describes the initial data collected from a wall outlet in Jacobs Hall on the UCSD campus. The data collected include voltage, frequency, and phase angles. We then describe the details of the CSSP system and how the data were collected. We also describe possible interpretations of the data as they relate to IEEE C37.118 and how the interpretations affect the data collected at UCSD [1].

## II. Frequency and Angle Measurement Data Collection

Angle measurements from PMUs are reported over the range of ±π radians or ±180 degrees. The rate of change of the angle is the frequency offset from the nominal frequency (60 or 50 Hz). Thus, if the system measured frequency is exactly 60 Hz, the reported phasor angle will appear stationary because the rate of change of the angle is zero. The frame of reference for the phasors is the 60 Hz nominal frequency, so the reference phasor rotates at this frequency and the phasor is measured relative to the nominal rotating phasor. The angle is measured relative to a standard bus operating at exactly 60 Hz with the peak of its cosine wave coincident with the top of the second wave. Thus the reported angle is relative to this virtual standard bus. The angle trend line shows an increasing or decreasing slope based on the offset from the nominal frequency.

Fig. 1 shows a 16-hour interval where a disturbance occurred at Jacobs Hall at UCSD. Data, such as voltage, frequency, and phase angle, were collected at a wall outlet in the hall and analyzed in order to locate a fault. The disturbance, which happened at about 6:09 a.m. on December 13, 2012, is visible near the middle of the trend.

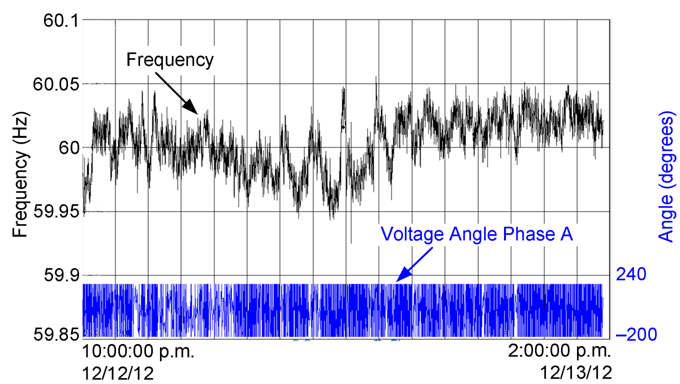


Fig. 1. Frequency and angle over a 16-hour interval

There are two scales in Fig. 1, one for frequency from 59.85 to 60.1 Hz and one for angle from –200 to 240 degrees. The blue line is angle, and the black line is frequency. The blue angle curve is shown displaced from the zero axis so the trend traces do not overlap. The frequency varies from less than 59.95 Hz to more than 60.05 Hz, or over a range of 55 mHz. What might appear as noise is actually how the grid behaves when responding to load and generation changes. The range of the angle is $\pm\pi$ radians, expressed in degrees on the chart. Note the rate of change of the angle measurement on the time scale in Fig. 1. It is so rapid that the trend chart almost looks like a solid blue line. Fig. 2 shows the same angle and frequency data as Fig. 1, but it is zoomed in to a time scale of 2 hours so that the disturbance is more visible.

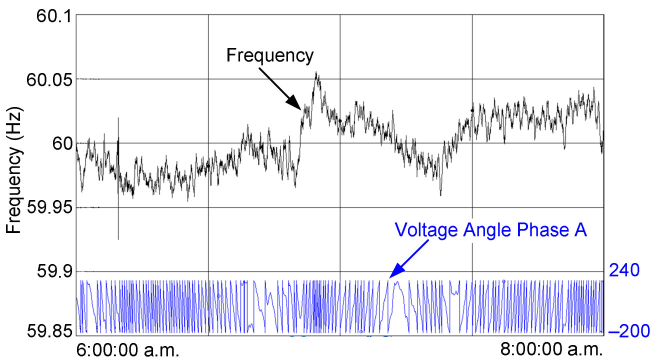


Fig. 2. Angle and frequency over a 2-hour period on December 13, 2012

Note that the slope of the angle trend in Fig. 2 indicates whether the frequency is above or below the nominal 60 Hz frequency. This view shows more clearly the rate of change of the angle measurement. It is wrapping in both the positive and negative directions based on the frequency difference from nominal. If the difference is positive, then the angle increases and the phasor rotates faster than the nominal in the counterclockwise direction. If the difference is negative, the phasor rotates in the clockwise direction.

Fig. 3 shows a more detailed view around the disturbance. A significant oscillation in the local grid, with the frequency dipping to less than 59.93 Hz, is clearly visible. This was likely caused by a sudden increase in load on the UCSD campus because there was no loss of local generation on campus. We can assume that this was a local disturbance because the Western Electricity Coordinating Council (WECC) Joint Synchronized Information Subcommittee (JSIS) has determined that wide-area events are defined as those having frequency changes of more than ±100 mHz.

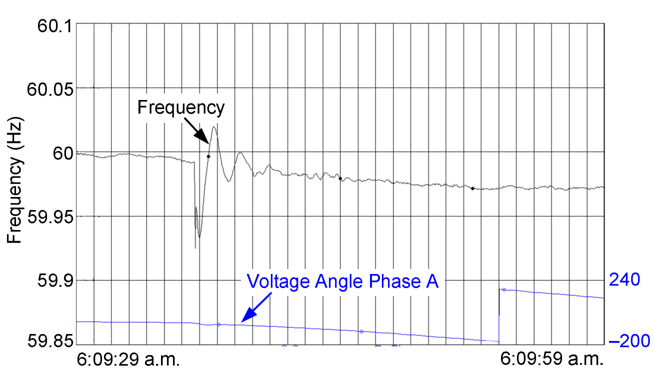


Fig. 3. Zoomed-in view of frequency and angle on December 13, 2012

A 6-second window displaying the oscillation in more detail is shown in Fig. 4.

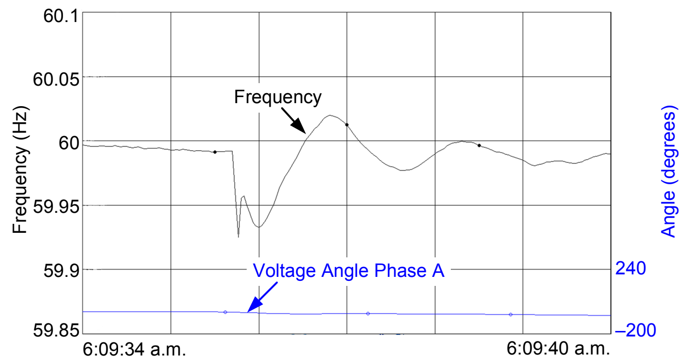


Fig. 4. A 6-second window showing December 13, 2012, event

The sudden drop in frequency followed by a successful recovery is clearly shown in Fig. 4. The start of the recovery is shown about 200 milliseconds after the nadir of the event and is followed by a second dip in frequency. This is a fairly common shape seen in the WECC grid during large generator-forced outages, and it has been observed repeatedly at UCSD.

We believe that the oscillations previously described are local to the UCSD campus, but system events occurring in the WECC grid are also observable in the 120 V single-phase measurement. For example, [2] describes the Chief Joseph brake test events, which can be measured by PMUs located in the WECC grid and at UCSD by 120 V single-phase PMU measurements. The Chief Joseph brake was built in 1975 to help characterize damping characteristics of the Western Interconnection. For testing purposes, 1,400 MW impulses are injected into the WECC grid at the Chief Joseph Substation in northern central Washington.

The UCSD PMU observations of two brake testing events on October 2, 2012, are shown in Fig. 5.

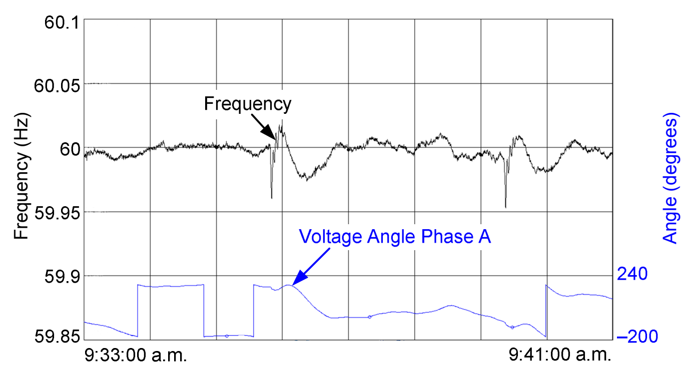


Fig. 5. Results at UCSD of Chief Joseph brake testing on October 2, 2012

The two Chief Joseph brake testing events were also recorded in Oakland, California. The frequency waveform is shown in Fig. 6.

We observed time delays of 261 milliseconds and 305 milliseconds between Oakland and UCSD, respectively, for these two events. For other events, such as generator trips, we have seen propagation delays of about 500 milliseconds. The frequency disturbance caused by the October 2 Chief Joseph brake testing event was first observed in Oakland and

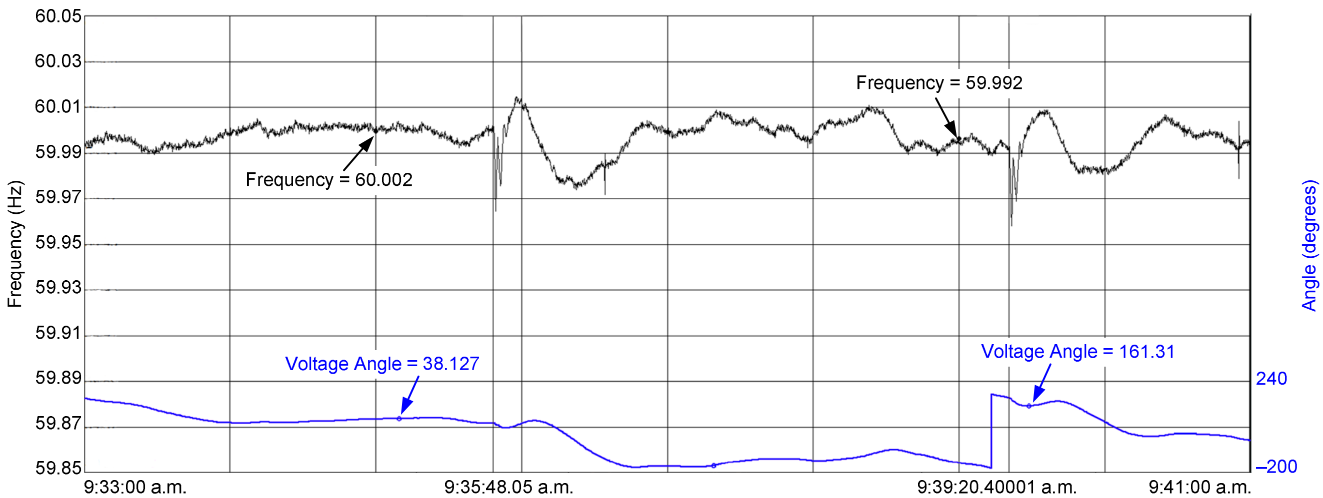


Fig. 6. Frequency and angle in Oakland for Chief Joseph brake tests on October 2, 2012

then, about one-third of a second later, in San Diego. Additional details of this analysis are available from UCSD upon request.

A discussion of the measurement system follows.

## III. Description of the Cybersecure Synchrophasor Platform

PMU data are now used for both wide-area monitoring, substation control and, in a few applications, for real-time feedback control. In order to provide a reliable signal for control, the measurements and the software interfaces should be redundant, as in some wide-area protection systems where the complete system is triply redundant. Reference [3] describes the importance of the redundancy and reliability of wide-area measurement synchrophasor historians and how they can be assembled using standard PMUs, standard commercial software, and industrial computers.

As described by the North American Electric Reliability Corporation (NERC), there is a need to architect and design systems that have commensurate levels of availability. NERC specifically states that synchrophasor system implementations should be done securely in redundant pairs to avoid systemic data gaps while standard maintenance is performed on the system [4].

The CSSP system that is currently in use at UCSD consists of two substation-hardened computers running a standard operating system and configured in a dual-redundant hot failover collective system, another substation-hardened computer running a standard operating system for local visualization, two PMUs, two switches, and a Global Positioning System-based (GPS-based) clock. The software includes a historian and software interfaces to read the IEEE C37.118 messages while computing FFTs, damping coefficients, and angle separations in real time. The CSSP software system is designed to handle up to ten PMUs per substation computer. Other devices in the substation that are transmitting standard protocols, such as DNP3, Modbus®, or OPC, can also be interfaced to the CSSP for improved availability.

The platform includes the following:

- High-availability historian systems running on two substation-hardened computers configured as a dual-redundant hot failover collective system.
- Real-time FFTs and real-time damping calculations running on each computer.
- IEEE C37.118 interfaces on each computer.
- A substation-hardened computer for local visualization.
- Redundant PMUs with dual GPS receivers.
- Redundant Ethernet switches.
- Visualization and application clients running on a third substation-hardened computer.
- A touchscreen display panel.
- A web portal and web displays.
- Software notifications configured for frequency exceedances.
- Equipment mounted in a single bay system.

The CSSP can accommodate up to ten PMUs sending data at 60 or 120 messages per second (mps) in IEEE C37.118 format. All PMU data are available on client software or compatible browsers inside the substation local-area network (LAN). PMU data (IEEE C37.118 packets) can be simultaneously and securely sent via multicast messages to other clients, including corporate host machines as well as regional reliability councils and independent system operators (ISOs). The system is extendable and fully programmable to perform real-time wide-area applications, including wide-area protection, control, and monitoring.

Data are locally archived in compressed form for 60 or more days, depending on the capacity of the disk drive installed. The oldest archive files are automatically stored offline using Microsoft® Background Intelligent Transfer Service (BITS) for long-term historical analysis. Compression settings can be used to reduce data volumes. The typical exception settings for a PMU are 50 percent of the accuracy specification for each measured value, and the typical compression setting is equal to the accuracy of each measured value. These settings are based on actual measured results that are described in [1] and are derived as follows. The recommended settings for exception are 0.5 times the reported accuracy of the device. If the PMU manufacturer specifies angle measurements to within 0.1 degrees, the exception setting would be 0.05 degrees. Similarly, the compression setting would be 0.05 degrees. With these settings, the loss of process information is greatly reduced. This is simply filtering out the random measurement error from the instrument. For example, the total vector error (TVE) requirements from IEEE C37.118 are 1 percent. So a conservative estimate of magnitude and angle exception and compression would be 0.5 percent for both angle and magnitude.

Additionally, the CSSP accepts data from Ethernet-connected devices in the substation that are associated with breakers, gas-insulated switchgear (GIS) switches, transformers, line monitoring devices, weather stations, and so on. The most common protocols for collecting these data are DNP3, Modbus, OPC, and Fast Message.

The architecture of the UCSD CSSP is shown in Fig. 7.

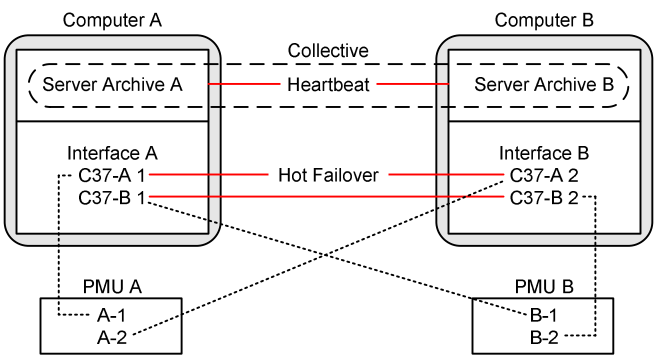


Fig. 7. CSSP architecture

Each rack-mounted PMU (i.e., PMU A and PMU B) sends data via IEEE C37.118 protocol messages to two separate IEEE C37.118 interfaces. Each interface (i.e., Interface A and Interface B) converts the IEEE C37.118 messages to individual data stream names and sends these names to a buffer subsystem that in turn sends the messages to the two servers (i.e., Archive A and Archive B), called a collective. No phasor data concentrator (PDC) function is required because when client software (like a display) requests data, the data are delivered as a time-synchronized data set; hence, there is no requirement for a PDC. Other devices in the substation can send data via Ethernet to the appropriate interface, and client software can also ask for time-synchronized data. In these cases, the device data that are received in forms other than IEEE C37.118 are interpolated to the IEEE C37.118 time stamps. This provides a time-synchronized data set for the application clients.

The interfaces are configured in hot failover mode. This means that the primary interface sends data to both servers while the secondary interface is reading data, configuring the data to be sent to the servers, and buffering the outputs in the event it becomes the primary interface. The data rates are typically 60 mps. The collective name is used when connecting to the server. The client software determines which machine to physically connect to depending on the load and the availability of the individual servers. The CSSP can be configured such that the local solid-state disk can hold 30 to 60 days of data for up to 40 PMUs and other devices in the substation. Compression settings can be applied to greatly extend archive history.

FFTs run in both computers and are executed at 60 mps data rates. The moving window is 512 points wide (8.35 seconds), providing a frequency resolution of 0.117 Hz. Time-synchronized angle differences of 60 Hz are also computed.

A large number of clients can connect to the data server and receive data formatted in OPC, web services, or OLE DB protocol.

## IV. Review of Additional UCSD Microgrid Oscillations

Several examples of oscillations inside the UCSD microgrid are shown in the figures in this section. Both frequency and angle are represented on the y axis, and time of day is represented on the x axis.

Fig. 8 shows an oscillation that occurred on November 3, 2012, that was likely due to a loss of generation, as indicated by the first swing in the negative direction.

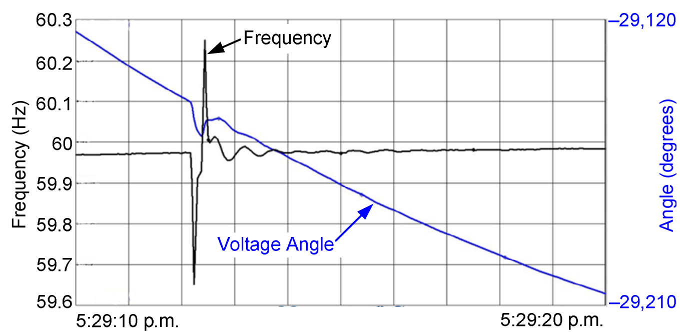


Fig. 8. UCSD microgrid oscillation example from November 3, 2012

As shown in Fig. 9, a significant oscillation in the UCSD microgrid occurred at about 10:23:10.5 p.m. on November 4, 2012. The frequency peak exceeded 60.12 Hz, and the duration of the oscillation lasted for more than 3 seconds. Note that the first swing is upward and was probably caused by a sudden loss of load. The angle measurement responded accordingly, showing a sudden increase at the time of the first swing. Recall that the frequency is the rate of change of the angle, so the PMUs are reporting accurately.

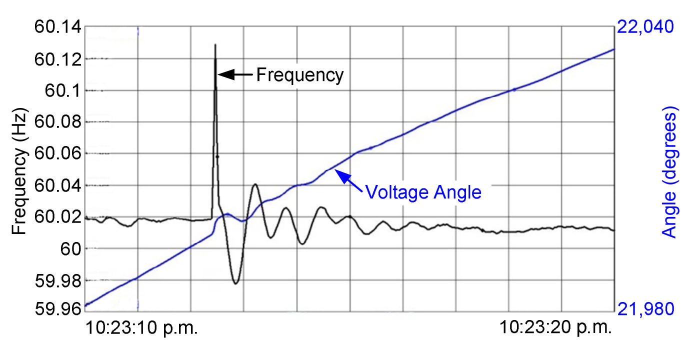


Fig. 9. UCSD microgrid oscillation example from November 4, 2012

A third oscillation example (see Fig. 10) shows similar behavior occurring on November 5, 2012, at about 2:42:56.8 a.m. This disturbance was also likely due to a load-loss event. This oscillation was also about 3 seconds in duration.

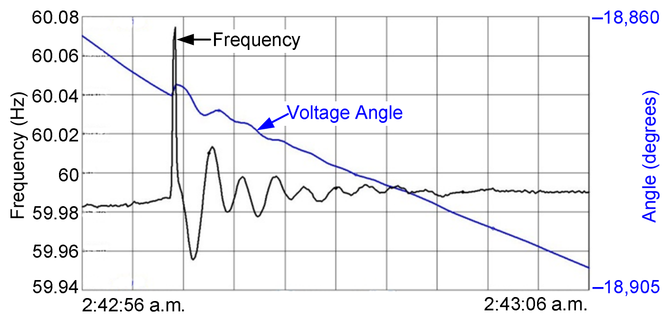


Fig. 10. UCSD microgrid oscillation example from November 5, 2012

## V. Next Steps

With the results of the initial installation and data collection at UCSD, the CSSP at UCSD will be upgraded to include newer versions of hardware and software, which will increase system availability, performance, and the storage of more archiving data. As shown in Fig. 11, the CSSP will be installed in the main switching East Campus 69 kV Substation, which feeds the entire UCSD campus, and will be connected at the 12 kV and the 69 kV windings of the three transformers feeding the campus from San Diego Gas and Electric. Fig. 12 shows the complete UCSD microgrid CSSP enterprise system after full installation.

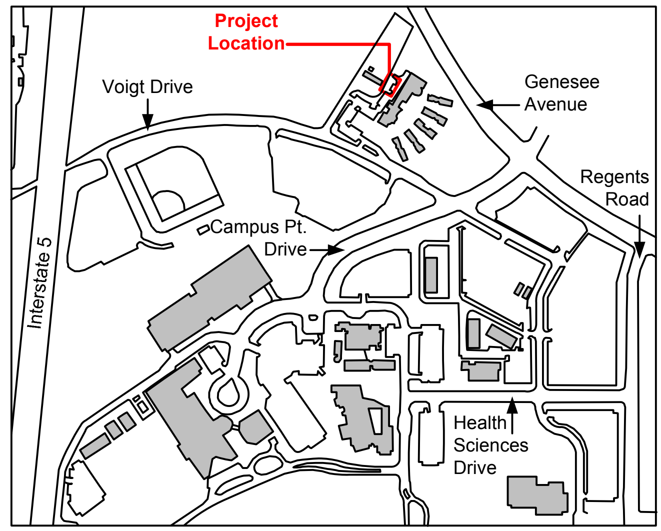


Fig. 11. East Campus 69 kV Substation map

With angle separations and FFTs of the data, the CSSP will then be used to provide advanced warning of stress developing in the area power system, allowing UCSD microgrid operators to prepare for a potential island maneuver.

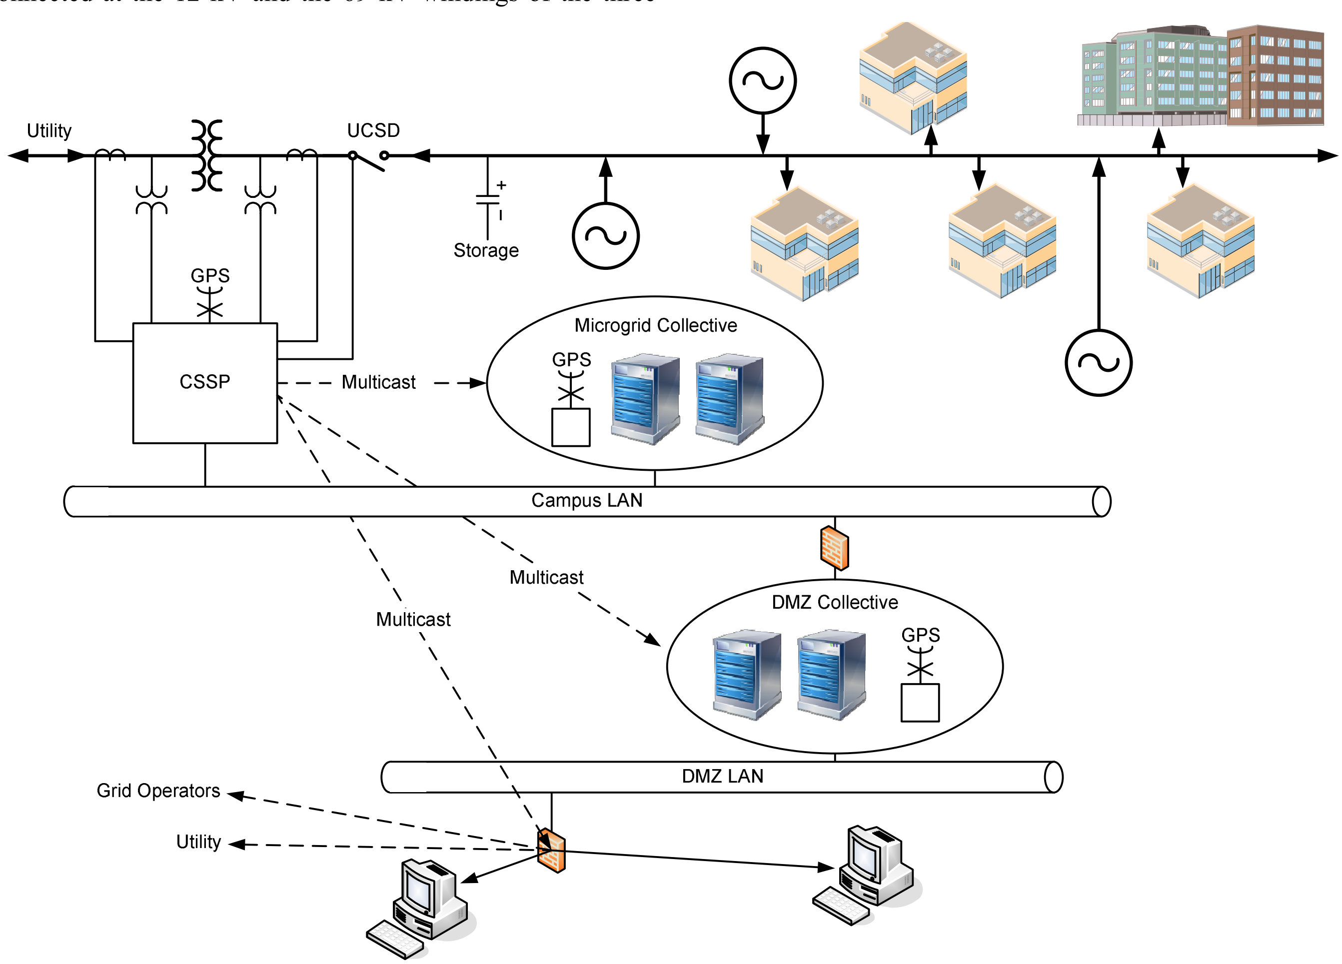


Fig. 12. UCSD microgrid CSSP enterprise system

## VI. Conclusion

The CSSP configuration is a convenient high-availability architecture for substation installation and the monitoring of distribution systems. The data are reliable and accurate and are used to compute real-time FFTs and angle differences.

## VII. Acknowledgments

This work was partially supported by the UCSD Smart Grid Investment Grant (SGIG) projects directed by Byron Washom, and funding was provided by the Jacobs School of Engineering intern program directed by Erica Kosa. UCSD students Daniel Paschall, Helen Zheng, and James Pitta were instrumental in configuring the CSSP computers, software, PMUs, and Internet connections.

## VIII. References

[1] IEEE Standard C37.118, IEEE Standard for Synchrophasor Measurements for Power Systems.

[2] M. Kintner-Meyer, R. Guttromson, D. Oedingen, and S. Lang, “Final Report for the Energy Efficient and Affordable Small Commercial and Residential Buildings Research Program – Project 3.3 – Smart Load Control and Grid Friendly Appliances,” Report No. PNNL-14342, Pacific Northwest National Laboratory, Richland, WA, July 2003. Available: http://www.pnl.gov/main/publications/external/technical_reports/PNNL-14342.pdf.

[3] C. H. Wells, J. Harrell, and R. Moxley, “Redundancy and Reliability of Wide-Area Measurement Synchrophasor Archivers,” proceedings of the 13th Annual Western Power Delivery Automation Conference, Spokane, WA, March 2011.

[4] North American Electric Reliability Corporation, “Real-Time Application of Synchrophasors for Improving Reliability,” October 2010. Available: http://www.nerc.com/docs/oc/rapirtf/RAPIR%20final%20101710.pdf.

## IX. Biographies

**Charles H. Wells** was born in Kokomo, Indiana. He graduated from Vanderbilt University with a B.S. in chemical engineering and from Washington University with an M.S. in chemical engineering and a Ph.D. 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.

**Gary Siefert**, P.E., E.E., is a member of the OSIsoft business development group, specializing in microgrid and renewable energy integration. He has been at OSIsoft since the fall of 2011 after a 31 year career in energy research at Idaho National Laboratory, a Department of Energy laboratory, Idaho Falls, Idaho. Mr. Siefert brings expertise and support to OSIsoft from his experiences as a senior program manager concentrating on renewable energy integration, resource assessment, SCADA cybersecurity, high-reliability power and control systems, military power system integration, IEEE 1547 standard development, and microgrid integration. He has also managed technical tasks for a multitude of federal and military agencies and installed total energy and water systems and incorporated renewable energy at high levels of penetration while preserving performance metrics. He has been a leader in the development of wind energy in Idaho, throughout the nation, and internationally and brings that expertise to the new smart grid integration of microgrids while preserving reliability and availability.

**William “Bill” Torre** is the program director for energy storage at the University of California, San Diego (UCSD) at the Center for Energy Research (CER) in the Department of Electrical and Computer Engineering. Mr. Torre recently retired as the manager of research and development at San Diego Gas & Electric (SDG&E) after working there for 30 years in various capacities, such as chief engineer, principal engineer, and manager of electric transmission engineering. Mr. Torre has led the development of smart grid projects at SDG&E, including the deployment of synchrophasors on the transmission and distribution systems and implementation of the Borrego Springs microgrid project, which included synchrophasors. He has also participated in several WECC committees that lead technical studies and the early testing of synchrophasors. He is a registered professional electrical engineer in California, holds bachelor’s and master’s degrees in electrical engineering, has published 14 technical papers, and holds one patent for a smart transformer.

**Byron Washom** is the director of strategic energy initiatives at the University of California, San Diego (UCSD), where he leads the university’s efforts to establish a highly innovative energy plan that will not only ensure that UCSD is able to meet both its present and future energy requirements in the most environmentally sustainable, cost-effective manner, but also serve as a premier example to institutions internationally. His previous achievements include eight technical world records for the conversion of sunlight to electricity while founder and CEO of Adcanco Corp.

Mr. Washom has been a key advisor on a wide range of clean energy technologies to the World Bank, the International Energy Agency, the U.S. Department of Energy, national laboratories, and numerous regulated utilities and corporations. He is a four-time Rockefeller Foundation grantee and a Heinz Foundation grantee. In 2008, Mr. Washom received the UCSD Community Leader Award for Sustainability and the R&D 100 Award for the 100 most outstanding innovations in the world.

**Ian Olson** works as a product engineer for Schweitzer Engineering Laboratories, Inc. (SEL) in the computing systems area. He joined SEL in 2005 to support customers with communications, protocols, and system integration. Now, as a product engineer, he manages the growth, direction, and technology of computer products. He graduated from the University of Idaho in 2004 with a bachelor’s degree in computer engineering.

**John Harrell** received his B.S. in mathematics at California Polytechnic State University in 1984 and his M.S. 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. His engineering work includes large parallel digital signal processor (DSP) firmware development and distributed control systems. He worked for Schweitzer Engineering Laboratories, Inc. from 2007 to 2013.


20130212 • TP6589-01