System for Event Summary Notifications: Preliminary Operational Results at SDG&E

Tariq Rahman
San Diego Gas & Electric

Ellery Blood, Matt Halladay, and Jared Bestebreur
Schweitzer Engineering Laboratories, Inc.

Presented at the
20th Annual Georgia Tech Fault and Disturbance Analysis Conference
Atlanta, Georgia
May 1–2, 2017

Originally presented at the
Power and Energy Automation Conference, March 2017

System for Event Summary Notifications: Preliminary Operational Results at SDG&E

Tariq Rahman, *San Diego Gas & Electric*
Ellery Blood, Matt Halladay, and Jared Bestebreur, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Modern power systems may include a variety of technologies configured to monitor electrical and physical parameters. Such monitoring systems provide a wealth of information regarding the health and operation of systems. Normally, these measurements are uniform and behave in a consistent manner through time. Areas of interest in these measurements, possibly indicating that a mechanical system is failing, are rare. The volume of these data can be overwhelming to operators attempting to identify these areas of interest. Automated identification of anomalous events offloads the job of continuous monitoring from the operators and allows them to focus their attention elsewhere.**

In this paper, we describe a frequency-domain statistical anomaly detector recently installed at San Diego Gas & Electric to collect streamed measurements and process them to produce event summary notifications. Simultaneous monitoring of multiple stations enables geographic, time, and frequency correlation of events. We further discuss results of operating this system at San Diego Gas & Electric.

I. Introduction

Phasor measurement units (PMUs) are an important tool for monitoring electrical transmission networks. They provide time-synchronized measurements (synchrophasors) continuously streamed at configurable rates through use of the IEEE C37.118 protocol. These measurements are aggregated at control centers and allow for convenient real-time analysis and visualization of systems. Visual analysis can identify many types of systems events, but effective real-time operation requires dedicated operation attention. Few automated analysis applications exist that are capable of performing such real-time analysis. This paper reports the algorithm and results of a real-time disturbance-detection system deployed at San Diego Gas & Electric (SDG&E).

SDG&E provides electricity to a 4,100-square-mile area of southern California. More than 100 PMUs provide continuous measurements at 30 samples/second. These data are used for real-time trending and historical analysis. This system generates more than three million data points per minute and more than six terabytes of data per year. With so many data points, manual analysis is impractical. An around-the-clock, automated disturbance-detection system was necessary for locating and classifying disturbances in synchrophasor data streams. Requirements for the algorithm include a low false-negative rate, algorithm tunability on a per-PMU basis, and simultaneous multi-frequency-range detection of oscillations.

A wide range of events occur in power systems, such as low-frequency inter-area oscillations, local power-electronic induced control oscillations, normal system reconfiguration events, and faults. The growing number of renewable energy sources connected to grids can also present a source of instability [1]. For example, SDG&E has observed power oscillations at approximately 10 Hz associated with a wind farm operating near peak capacity.

Several synchrophasor-based wide-area measurement systems (WAMS) have been previously described in [2], [3], [4], [5], and [6]. Of specific relevance to our algorithms and methods are the publications from Bonneville Power Authority (BPA) regarding their Oscillation Detection Monitor [5] and the algorithmic research from the University of Texas for event detection in the Texas Synchrophasor Network [1]. The purpose of this publication is to supplement this body of knowledge with real-world operational results of such a system and to share specific adjustments that have been employed to improve algorithm performance.

II. Algorithm

A. Methodology

When choosing an algorithm to assist power system operators to detect oscillatory events, we first have to reiterate that the human in the loop, either operator or engineer, is the ultimate decision maker regarding event detection and classification. A detection system does not inherently need to characterize and catalog events; it can provide value by simply informing the operator that something unusual is occurring and present them with appropriate information with which to make a decision. We therefore focused first on building a system that was effective at detecting anomalous behavior with respect to oscillatory events.

First, we collect short windows of power system data and convert them to the frequency domain via a discrete Fourier transform. Basic preprocessing, such as mean removal and windowing, is applied to reduce signal-processing artifacts.

Next, we sum the oscillatory energy contained in various predefined frequency ranges. Several ranges have been proposed in previous literature [7] [8]. In our initial implementation, we chose frequency bands proposed by BPA for monitoring oscillations on the Western Interconnection [7] because SDG&E's system is also located in this region. These frequency bands are 0.01–0.15 Hz, 0.15–1.0 Hz, 1.0–5.0 Hz, and 5.0–14.0 Hz.

The values for the respective band oscillatory energy are recorded and indexed for quick searching and retrieval. On their own, these values are suitable for identifying potentially damaging oscillations if expert knowledge is available on equipment susceptibility to power fluctuations. However, our goal was to develop a system requiring minimal preconfiguration.

Because the oscillatory energy by itself does not provide an indication as to the presence of anomalous oscillatory behavior, we then looked to statistical analysis methods. One prevalent method for identifying anomalous values in a data set involves use of the z-ratio. The z-ratio [9] is a statistic calculated by taking a sample value, x, and normalizing it by subtracting the mean and dividing by the standard deviation, as shown in (1). This results in a unitless value that represents by how many standard deviations a sample value deviates from the mean. In practical applications, the mean and standard deviation values are calculated from a sliding one-hour window of data preceding the present measurements. These z-ratios are similarly recorded and indexed.

$$z = \frac{x - \overline{x}_{previous\_hour}}{\sigma_{previous\_hour}} \tag{1}$$

B. Implementation Details

1) Statistics

When calculating the mean and standard deviation for z-ratio calculations, it is important to remember that the power system is constantly evolving. As such, running recursively updated statistics for weeks or months would tend to produce false positives because the behavior on a summer day might be highly anomalous when compared with statistics recorded during the winter. We found that calculating the statistics over a rolling one-hour window immediately preceding the present values provided good results.

We did, however, identify the need to temporarily suspend statistical updates during an event. Because the z-ratio during an event could be in the hundreds of standard deviations, updating the statistics during the event resulted in a large statistical skew. This skew resulted in two undesirable operating characteristics. First, detection of a sustained oscillation tended to conclude before the actual event concluded. As the anomalous behavior worked its way into the statistics, that behavior was no longer considered anomalous. Second, after the conclusion of an event, the statistics were artificially high, making the z-ratio less sensitive to subsequent events. This could potentially mask new events for as long as an hour after the first event concluded. Suspending the update of statistics during an event effectively removed both of these undesirable artifacts.

2) Detection Thresholds

Identification of an oscillatory event is accomplished through the comparison of the present z-ratio to a predefined threshold. From analysis of data provided by SDG&E, we found that an instantaneous pickup with a threshold of thirty standard deviations worked well. Thus, we would alert operators if a measurement of oscillatory energy exceeded the previous hour's mean by 30 times the previous hour's standard deviation. We also identified the need to apply alternative thresholds on a station-by-station basis. Some stations, due to proximity to noise-inducing equipment, would fail to detect an event. If future analysis of operational data reveals such a situation, we will set the threshold for PMUs at that station to a lower level.

We also found that many momentary conditions created spikes in the z-ratio but were uninteresting to operators. To filter these spurious events out, we included a pickup timer so that an event notification would only be triggered if the z-ratio exceeded its threshold for a user-specified minimum duration. We intend to fine-tune the settings based on operator feedback and review of operational data. We anticipate that this fine-tuning will result in a smaller threshold for disturbance level with a non-zero pickup time to more effectively identify real events while filtering out spurious data.

3) Simultaneous Events

When an oscillatory event of interest occurs, it is often detected simultaneously by multiple stations. In order to not overload operators with redundant information, we designed the system to automatically group similar events that overlap chronologically into one event notification.

Once these events are consolidated, we compare the relative magnitudes of the energy and z-ratio across the stations. The list of affected stations is then sorted by disturbance level; highest on top, with less-strongly affected stations below. This aids operators in identifying the location of the event because the stations most closely coupled with the event will often have similar energy values, and a sharp drop in disturbance level will indicate which stations are less affected.

4) Threshold and Pickup Tuning

In addition to calculating statistics and z-ratios on newly arriving data, this system additionally calculates these values for the historical archive. The indexed energy and z-ratios (Fig. 1) can be easily searched to identify suitable thresholds that will indicate historical events of interest while minimizing spurious alerts. Here we considered the tradeoff between the nuisance of false positives (spurious alerts) resulting from a low threshold versus the operational impact of false negatives from a high threshold. A false positive results in a nuisance alert that costs operators a few seconds to review and clear. False negatives can potentially result in system damage due to degraded situational awareness and delayed remedial action. Given this tradeoff, we chose to configure a low threshold to minimize the probability of false negatives.

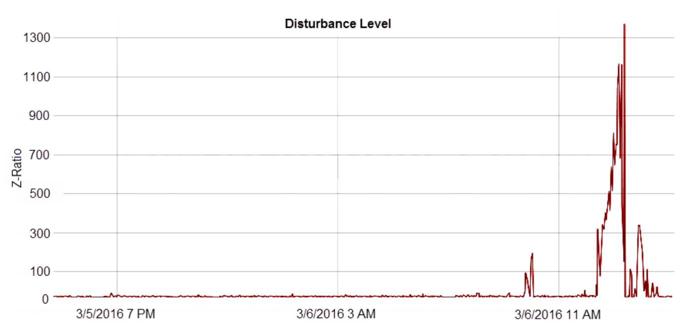

Fig. 1. Software display from detected oscillation

With this constraint in mind, we propose the following method by which to tune the detection parameters. First, an engineer would employ the built-in search function to find the largest z-ratio threshold that detected all the events of interest. Once a suitable threshold is identified, gradually increase the pickup time to reduce the number of spurious detections while retaining all desired detections.

III. San Diego Gas & Electric Results

A. Introduction to SDG&E Disturbance Results

The authors developed and deployed an automated disturbance-detection algorithm into SDG&E's wide-area situational awareness software application. Due to the large number of PMUs that SDG&E has deployed and the limited availability of operators and engineers to analyze these data, automatic detection of disturbances and oscillations is of significant benefit to SDG&E.

The developed disturbance-detection algorithm is currently deployed at SDG&E. The results shared in this paper represent a summary from one week of data gathered at SDG&E. Fig. 2 summarizes the daily count of detected disturbance per day for this period.

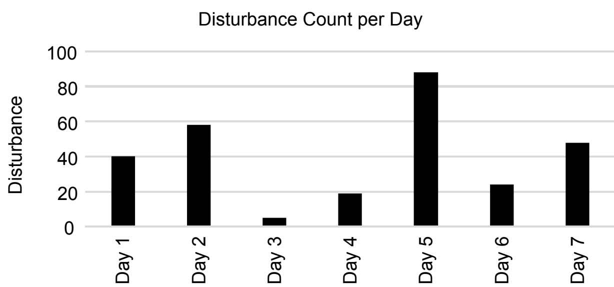

Fig. 2. Distribution of detected disturbances over a one-week period

The detected disturbance count varies greatly day to day. Many of the detected disturbances are likely linked to the same root cause. As the severity of the resulting oscillations varies above and below the detection threshold, multiple related event notifications may be generated. A significant number of the detected disturbances for Days 1, 2, 5, and 7 (as shown in Fig. 2) were tied to repeated oscillations attributed to the same root cause. This is the primary reason for the spike in disturbances when compared to Days 3, 4, and 6.

To effectively diagnose and mitigate the causes of these oscillations, operators and engineers need to be able to quickly identify, locate, and analyze them for impact. To facilitate this, the oscillation detection software includes an event summary notification to quickly alert operators to the situation.

Fig. 3 through Fig. 5 show examples of the event summary notification displayed by the software. This is the primary display by which an SDG&E operator or engineer interfaces with the oscillation detection system. The left side of the notification displays plots of system frequency (top), voltage magnitude (middle), and disturbance level (bottom) for the five most strongly affected PMUs. The system frequency plot shows the raw data from which the disturbance level is calculated. The disturbance level plot represents the z-ratio as described in (1) and provides a general indication of how abnormal the present condition is. The voltage magnitude plot is also included as some detected events are more visible to operators via the voltage magnitude than frequency signals.

The right side of the event notification provides quantitative information regarding to what extent various PMUs were affected by the event.

The notification displays the 15 most strongly affected PMUs ordered by disturbance level. To aid in comparison, a percentage of maximum disturbance level is also provided. For example, the top station always has 100%; subsequent stations show their disturbance level divided by the top stations disturbance level expressed as a percentage. Below this list, the notification indicates how many stations in total detected this event.

Also on the right side, the notification identifies in which of the four frequency band(s) the event was detected and the frequency of peak energy.

B. Example 1–Local Oscillation

The increase of renewable energy penetration in the form of solar and wind generation on SDG&E's system is changing the operation of the grid and introducing new challenges in the form of high-frequency oscillations.

In the example shown in Fig. 3, the system identified an oscillation in the 5–14 Hz frequency band associated with three PMUs whenever the solar plant reached peak power.

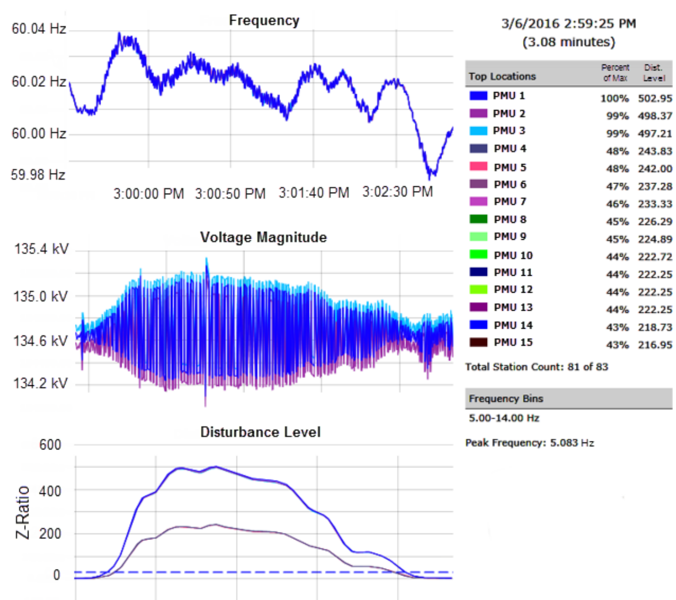

Fig. 3. Software display from detected oscillation

In this example, the high-frequency oscillation is clearly visible to the operator via the voltage magnitude plot (Fig. 3, center). PMU 1, PMU 2, and PMU 3 each show a percentage at or above 99%, with the next largest below 50%. These three PMUs are all located in the same substation. PMU 4 and PMU 5, with impact levels approximately 50% of PMU 1, are located in neighboring substations. The oscillation originated from a solar farm connected to the substation being monitored by PMU 1, PMU 2, and PMU 3. This particular oscillation was detected by PMUs throughout SDG&E's system.

The 5.08 Hz disturbance was grouped into the 5–14 Hz frequency band typically associated with generators, HVDC, and SVC [7]. This encourages the operator to investigate high-frequency sources like solar or wind farms near the substation containing PMU 1, PMU 2, and PMU 3.

C. Example 2–Line Trip Across Utility Interconnection

In addition to detecting abnormal power system conditions such as oscillations, synchrophasor data also provide operators and engineers additional insight and wide-area context for traditional power system operations such as line and generation trips. Due to the real-time nature of synchrophasors, Operations personnel at SDG&E often use synchrophasor data to determine faulted phases for a line trip prior to collection and analysis of the relay event records. The real-time synchrophasor data also provide visibility into the reclosing of transmission lines (in particular, inability to close due to too large of an angle difference across the line). The process of analyzing these line trips is greatly simplified by automatic detection and operator notification.

Fig. 4 shows an example display of a line trip detected on a line between SDG&E and a neighboring utility. PMU 1 and PMU 2 monitor the 230 kV bus and lines across the interconnection. The disturbance impact at these PMUs is significant when compared to the rest of the system, as is typical for line trip events. PMU 3 is monitoring a 138 kV line out of the substation, and PMU 4 and PMU 5 monitor two 500 kV lines out of the substation. The remaining PMUs are located in neighboring substations within SDG&E's system.

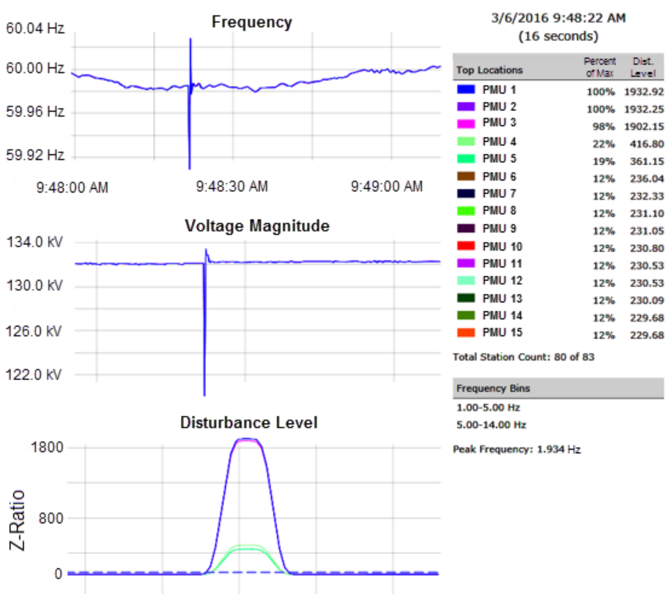

Fig. 4. Software display from detected line trip

D. Example 3–WECC Generation Trip

The goals for utilities like SDG&E that cover relatively small geographic areas when compared to the size of the Western Interconnection are different when compared with those of transmission system operators (which are the traditional consumers of disturbance and oscillation detection tools). SDG&E Operations personnel are primarily interested in detecting disturbances and oscillations that occur within or near their geographic region. If the disturbance occurs within their region, they are likely able to take corrective action and improve the reliability of the power system for their customers.

The disturbance-detection algorithm deployed at SDG&E is currently based on detecting statistical anomalies in frequency measurements. As a result, disturbances and oscillations that occur throughout the entire interconnection will be detected. Because SDG&E Operations is primarily interested in disturbances within their region, it is important for an operator or engineer to be able to quickly determine if the disturbance source location is within their region or not.

Fig. 5 shows an example of a generation trip that occurred outside of SDG&E's region of the Western Interconnection. In general, SDG&E operators and engineers can determine if a disturbance is out of their system by looking for the following criteria in the software display:

- The disturbance impact for the top 15 PMUs are within several percentage points of one another.
- Every PMU within the system detects the disturbance.

This method of determining whether a disturbance is in-system or out-of-system assumes that the electrical distance between SDG&E's PMUs is negligible when compared to the electrical distance to disturbance location. This method of determination may not work for disturbances that occur in neighboring systems, but it can be augmented by incorporating data from several PMUs in the neighboring systems.

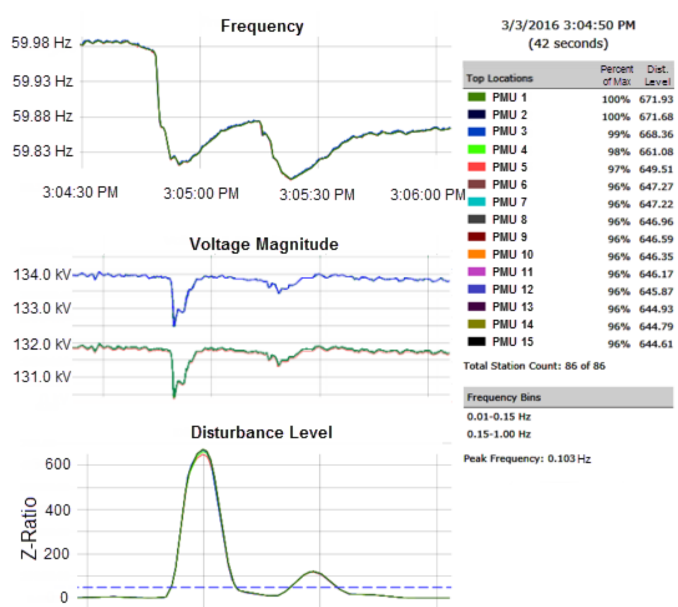

Fig. 5. Software display from detected WECC generation trip

IV. Conclusion and Future Work

This system has been successfully deployed at SDG&E. The automatic detection and notification capability has freed SDG&E operators and engineers from the burden of continuously monitoring hundreds of PMU streams to manually identify events. As more PMUs are deployed to SDG&E's system, especially in distribution, this capability will be even more valuable to SDG&E operators and engineers.

In addition, the system has detected a few anomalies that were not visible to the eye from directly viewing the PMU stream. These are often small-magnitude, high-frequency events that are typically hidden in the noise. They are detected when a PMU that typically has minimal energy in the 5–14 Hz band begins to exhibit oscillations in this range. We anticipate that detailed study of these events may provide additional insight into control system calibration drift or impending equipment failure.

V. Acknowledgements

The authors would like to acknowledge the support of Subbu Sankaran (Software Component Architect), Dan Eklund (Retired from San Diego Gas & Electric), and Dan Dietmeyer (Senior Operations Engineer at San Diego Gas & Electric) for their contributions to the development and deployment of the system within SDG&E's engineering and operations environments.

VI. References

[1] S. W. Sohn, A. J. Allen, S. Kulkarni, W. M. Grady, and S. Santoso, "Event Detection Method for the PMU's Synchrophasor Data," *IEEE Power Electronics and Machines in Wind Applications (PEMWA)*, Denver, July 2012.

[2] P. Mukhopadhyay, R Nagaraja, R. Anumasula, F. Khan, C. Kumar, and D. N. Kumar, "Disturbance Analysis Tool Based on Synchrophasor Data," *2014 IEEE International Conference on Power Electronics, Drives and Energy Systems (PEDES)*, pp. 1–6, 2014.

[3] Y. Ota, "Disturbance Detection of Wide Area Power Systems by Use of Short-Term Synchrophasor Data," *2015 IEEE Eindhoven Power Tech*, pp. 1–4, 2015.

[4] D. Zhou, J. Guo, Y. Zhang, J. Chai, H. Liu, Yong Liu, C. Huang, X. Gui, Yilu Liu, "Distributed Data Analytics Platform for Wide-Area Synchrophasor Measurement Systems," *IEEE Transactions on Smart* Grid, vol. 7, no. 5, pp. 2397–2405, Sept. 2016.

[5] A. Donahoo, "BPA's Oscillation Detection Monitor," presented at the Joint Synchronized Information Subcommittee (JSIS) Meeting, Salt Lake City, Utah, 2016.

[6] G. Liu and V. Venkatasubramanian, "Oscillation Monitoring from Ambient PMU Measurements by Frequency Domain Decomposition," *proceedings of the 2008 IEEE International Symposium on Circuits and Systems*, pp. 2821–2842.

[7] N. Leitschuh, "Synchrophasor based Oscillation Detection at Bonneville Power Administration," presented at North American SynchroPhasor Initiative (NASPI) Meeting, 2014.

[8] G. Rogers, *Power System Oscillations*, New York: Springer, 2000.

[9] E. Kreyszig, *Advanced Engineering Mathematics*, 4th ed., New York: Wiley, p. 880, 1979.

VII. Biographies

Tariq Rahman is a principal engineer in the Electric Transmission and Distribution Engineering Department at San Diego Gas & Electric (SDG&E). Prior to joining SDG&E, Tariq worked with Long Island Lighting Company in New York and the Sunflower Electric Power Corporation in Garden City, Kansas. He obtained his B.S.E.E from Bangladesh University of Engineering and Technology in 1980 and his M.S.E.E from the University of Arkansas, Fayetteville, in 1985. Tariq has 30 years of experience working in the electric utility industry. During this time, he has worked in electric operations, generation planning, and system protection and control engineering. In addition to his system protection engineering functions, he is leading the Transmission Synchrophasor Project at SDG&E. He is a licensed professional engineer in the states of New York and California and is a member of IEEE.

Ellery Blood, Ph.D. is a senior research engineer at Schweitzer Engineering Laboratories, Inc.; he provides strategic support of synchrophasor and event analysis software products as well as research into wide-area control. Before coming to SEL, he taught Systems Engineering as an assistant professor at the United States Naval Academy and is an active Naval Reservist supporting the Office of Naval Research where he serves as the National Director of Reserve Support of Education Outreach. He holds both a masters and doctorate degree in electrical and computer engineering from Carnegie Mellon University. He also earned an M.S. degree in mechanical engineering from the Naval Postgraduate School and a B.S. degree in computer and systems engineering from Rensselaer Polytechnic Institute. He is an active member in the North American SynchroPhasor Initiative (NASPI) community.

Matt Halladay is a development lead engineer at Schweitzer Engineering Laboratories, Inc.; he is a developer in the analytics software group and provides support for both research and algorithms development in the areas of digital signals analysis, compression, and visualization. He holds a B.S. degree in Computer Science from Troy University and also earned an Associate of Arts degree in Arabic Language Studies form the Defense Language Institute.

Jared Bestebreur is a product manager at Schweitzer Engineering Laboratories, Inc.; he is responsible for synchrophasor and relay event analysis software products. He has supported the implementation of wide-area monitoring systems around the world. Jared holds a B.S. degree in electrical engineering from Washington State University. He is also an active member in the North American SynchroPhasor Initiative (NASPI) community.

20170127 • TP6792-01