

RIT/PSY/TR-18/1
Final Technical Report

Improving Operator Situation Awareness by Phasor Measurement Unit (PMU) Data Visualization

By

Esa Rantanen, Ernest Fokoué, Jacob Haut, and Chandini Ramesh Rochester
Institute of Technology, Rochester, NY

Kathleen Gegner
University of Illinois at Urbana-Champaign

Thomas Overbye and Komal Shetye
Texas A & M University, College Station, TX

For

Bonneville Power Administration, Technology Innovation — ST3
905 NE 11th Avenue, Portland, OR 97232

May 1, 2018

Technical Report Documentation Page

1. Report Title Improving Operator Situation Awareness by Phasor Measurement Unit (PMU) Data Visualization	2. Report No. RIT/PSY/TR-18/1
	3. Report Date May 1, 2018
4. Authors Esa M. Rantanen, Ernest Fokoué, Jacob Haut, Chandini Ramesh, Kathleen Gegner, Thomas Overbye and Komal Shetye	5. Type of Report Final Technical Report
6. Performing Organization Name and Address Rochester Institute of Technology, Rochester NY University of Illinois at Urbana-Champaign, IL Texas A&M University, College Station, TX	7. Performing Agency Code NA
8. Sponsoring Organization Name and Address Bonneville Power Administration, Technology Innovation-ST3 905 NE 11th Avenue, Portland, OR 97232	9. Sponsoring Agency Code NA
10. Supplementary Notes This research was supported by a grant from the Bonneville Power Authority (BPA) TIP#353: "Improving Operator Situation Awareness by Phasor Measurement Unit (PMU) Data Visualization". The BPA project manager was Mr. Gordon Matthews, ST3.	
11. Abstract The application of phasor measurement unit (PMU) data in the power industry is currently an area of intense interest. The key driver for PMU technology is to use the precise time sources provided by Global Positioning System (GPS) satellites to accurately measure the relative voltage and current phase angles at buses across an interconnect at a very high sampling rate. Presenting PMU data to power system operators in a format that is truly useful for them and that affords improved situation awareness (SA) and fast and accurate decision making is a particular challenge to display design. This report describes development of prototype displays of PMU data where the design criteria were derived from characteristics of the data as well as the situation awareness requirements of power system dispatchers. A prototype display that meets all these design criteria is described.	

TIP#353 Final Technical Report:

Improving Operator Situation Awareness by Phasor Measurement Unit (PMU) Data Visualization

Prepared For

Bonneville Power Administration, Technology Innovation — ST3
905 NE 11th Avenue, Portland, OR 97232

Prepared By

Esa Rantanen, Ernest Fokoué, Jacob Haut, and Chandini Ramesh
Rochester Institute of Technology, Rochester, NY

Kathleen Gegner
University of Illinois at Urbana-Champaign

Thomas Overbye and Komal Shetye
Texas A & M University, College Station, TX

May 1, 2018

Executive Summary

Operators in electric power system control centers serve a critical role in ensuring the integrity of the nation's electric grid and to prevent disruptions (e.g., blackouts) to the supply of electric power at any time, for any reason. The enormity and complexity of the grid itself is mind-boggling. However, recent advances in both display technology and measurement of various aspects of the grid have made display of truly dynamic system data both possible and affordable.

The application of phasor measurement unit (PMU) data in the power industry is currently an area of intense interest. The key driver for PMU technology is to use the precise time sources provided by Global Positioning System (GPS) satellites to accurately measure the relative voltage and current phase angles at buses across an interconnect. Presenting PMU data to power system operators in a format that is truly useful for them and that affords improved situation awareness (SA) and fast and accurate decision making is a different problem, however.

The main objectives of our research were to (1) perform a Cognitive Work Analysis (CWA) on power system operators to identify their information needs and uses, (2) develop visual displays of PMU data that directly support operators' SA and decision making, (3) validate the displays to demonstrate improvement in operator performance with the aid of the new PMU data displays, and (4) provide BPA with human factors guidelines for evaluation of all future visualization solutions.

The results of the CWA are given in Part I of this report. The key to a successful CWA is engagement with the subject matter experts (SMEs) in the work domain of interest. We had several very valuable opportunities to observe the dispatchers at the Dittmer control center (DCC) and speak with them about their work, different tasks they do, and the tools at their disposal. These observations provided a good foundation for the CWA. However, subsequent attempts to "fill in the blanks" of the CWA were hampered by our inability to collect more detailed data from the DCC dispatchers. Only 11 dispatchers responded to an online survey and their responses were too varied for clear conclusions. Nevertheless, the CWA afforded several key insights for our subsequent efforts in PMU data visualization. The primary insight was that the PMU visualizations should allow the dispatchers move away from the current alarm-driven, reactive, mode of operations to analytical, proactive, mode. Second, we designed our visualizations to be displayed on the video wall so that the information they provide would be shared by all dispatchers at all times.

We developed a total of 5 different visualizations of PMU data. In Part II of this report we describe a prototype display consisting of 6 panes, showing three key variables (voltage, frequency, phase angles) at the same in three different formats (underlying grouping of data based on the maximum differences between time steps for each bus in histograms, line plots against time of the aforementioned variables for the outlier buses, and geographic locations of the buses on a map). This display captures the physical footprint of the test system as well as the temporal component of the measurement data.

In Part III of the report we describe a prototype with three key design features: (1) visual isolation of "misbehaving" buses by displaying first-order differences or rates of change, (2) presentation of the three most critical variables on a common time axis so that their interactions may be easily perceived, and (3) "compression" of the reverse logarithmic time axis so that as much history (to the left of the axis) as possible may be viewed at all times while most recent events to the right of the axis are displayed with high resolution.

The other two prototype displays described in Part III are geographic. One is a "heat map", or a color-

contoured map based on a mix of univariate nonparametric novelty detection and parametric outlier detection. The heat map generation does not look at the substations, but instead focuses entirely on each bus as its own entity. The heat map is generated by extracting geographic coordinates of each bus along with its contours, constructing a polygon to represent the bus region on the map, and the values of the rates of change in the variables are color-coded. We also explored various clustering methods to identify anomalies and outliers in the three key variables measured at each PMU-equipped bus in the system. These anomalies were displayed as pins on a geographic map. Note that time in map-based display must be represented dynamically by animation (i.e., the color contours change and shift, or the pins appear and disappear in time as the grid event unfolds). Different events to produce different dynamic patterns that may be recognizable by experienced operators. Although movement in animated displays can be a very powerful attention-getter, pattern recognition comes at the cost of watching the display for some time. In other words, an animated map-based display cannot be read “with a single glance” like the stacked line-plot display could.

Finally, Part IV of this report we present novel visualizations of power system modes and oscillations, to enable better situational awareness and decision making. The displays are color-contoured maps, with associated line-plots against time axis.

We did not have the opportunity to test these visualization solutions created with this project in an operational setting, or even simulations with SMEs. Therefore, they are offered as prototypes that should be subjected to rigorous evaluation and analyses of their utility in power system control. We also did not manage to create just one display that would have met all the criteria we had generated. At a minimum, a two-panel display seems necessary, where a geographic map view is paired with a display where the data are plotted against a time axis.

Part I

Preliminary Research and Simulations

1 Background

In 2003, the National Academy of Engineering recognized electrification in the United States—and the infrastructure that made it possible—as the greatest achievement of the 20th century [1]. While still remarkable, the U.S. electric grid of the 21st century looks almost exactly like that of the 20th century, and though the electricity still flows through the aging system, it is doing so less reliably and efficiently. In fact, among all developed nations, the United States has the highest number of outage minutes, which not only leads to annoyed customers, but also economic loss for businesses and utilities [2]. Less reliable service is not the only shortcoming of the aged 20th century electric grid. Cyber and physical attacks, the environmental impact of energy production, updating the digital and communication network, dealing with more extreme weather, and leveraging data for flexibility and controllability, are additional growing pains the future grid needs to address.

The future electric grid, or the so-called “smart grid” refers to an electric grid that is able to integrate a large amount of renewable energy sources, despite their intermittency, and has the latest and most secure communication and digital infrastructure available. It also should handle the new paradigm of distributed generation (e.g., microgrids) rather than centralized generation (i.e., macrogrid) and be more flexible and able to adapt to changes in generation and load more quickly than the present system. Quick recovery from disturbances is a key feature of the smart grid, requiring provision of better more information faster to operators at “smart” control centers [3, 4] (see also [5, 6, 7]).

The intelligence of the future electric grid, or data and how they are used to improve operation of the electric grid, is most relevant to this research. Grid intelligence requires that information about the state of the electric grid be collected and shared. For the 21st century, the added requirement for grid intelligence is that more information be shared more quickly, and more securely than before. To collect more information, more sensors have been (or are being) deployed to collect measurements throughout the U.S. electric system. At the transmission system level, these sensors are called synchrophasors or phasor measurement units (PMUs). Similar sensors, like smart meters, are being deployed to homes at the distribution level. Additionally, the grid’s communication infrastructure, both at the transmission and distribution levels, is being improved to ensure that all the new sensor data can reach its end user quickly and safely.

With many sensors already installed and a communication infrastructure mostly in place, the next step is what is being done with the data that is collected by all the new sensors. To truly claim intelligence, it is not enough for the grid just to collect and share data, but the data must enable new insights and drive better decisions. Utilities, whether at the transmission or distribution level, have access to much data, but lack the means to turn it into something really meaningful. The problem has therefore shifted from scarcity of relevant data to overabundance of them and ultimately to human limitations in processing all the available information. Human judgment and decision making are still critical in operation of the electric grid, despite advances in automation and computers now performing many routine tasks. The role of a human operator is particularly critical during unforeseen (unforeseeable) events that computers have not been programmed to handle [8] [9] [10] [11].

1.1 Grid Operation and PMUs

In the transmission system, operation of the grid requires maintaining the system frequency to be near 60 Hz, balancing tie-line flows (transmission lines that carry electricity between neighboring utilities), and ensuring line currents, equipment loading, and voltage levels are within their limits [12]. Duties of, and decisions made by, operators largely revolve around the above. Thus, to improve operators' performance, data must provide keener insights about one or more of those operational objectives.

Historically, operators have relied on information gleaned from supervisory control and data acquisition (SCADA) systems, which has been sufficient for normal, localized operation of the grid. But, when trouble arises, more precise, frequent, time synchronized, and widespread measurements can help. Unlike SCADA measurements, which are collected every 2 to 4 seconds, PMUs record data at a minimum of 30 times per second (i.e., at 30 Hz), with each measurement time stamped using the global positioning system (GPS). These data include voltage magnitude and angle, current magnitude and angle, and frequency, measured at the buses (i.e., conductive connection points between electrical equipment) where they are installed. The precise GPS timestamp allows data from many PMUs to be integrated to provide comprehensive visibility of an electric grid and its neighboring systems. That widespread visibility is referred to as wide-area situational awareness (SA), and it is a key reason for using PMUs as it allows operators to see problems as they develop in neighboring systems, before they reach their own. In fact, one of the main reasons for the Northeast Blackout of 2003 was that utilities did not have enough visibility and knowledge of their neighboring utilities' systems [Reference needed]. PMUs circumvent that problem by delivering data to all parties who want it.

Furthermore, given the time scales for various power system events, and the monitoring capabilities of PMUs and SCADA systems, PMUs capture events such as switching surges, inverter-based controls, and stator transients and subsynchronous resonance. These are outside the seconds-scale resolution of SCADA but within the 10^{-3} s resolution of PMUs [13].

PMUs are installed at substations and record measurements for multiple buses within a substation. A typical PMU installation connects potential and current transformers to various buses in the substation to record voltage and current measurements. A single PMU has multiple channels through which measurements can be recorded, allowing one PMU to collect measurements for multiple buses. Once measurements are made by a PMU, their data is sent to a regional hub, called a phasor data concentrator (PDC). PDCs time-synchronize and aggregate data from many PMUs, using the timestamp assigned to the data by the PMU's GPS clock. Once aggregated, the data is then sent to another PDC, for further aggregation, or to a central control room for an operator to deal with.

As of 2015, more than 1,500 networked PMUs were installed in North America, and more than 225 PDCs, as a result of funding from the American Recovery and Reinvestment Act of 2009. This larger PMU infrastructure has given operators visibility of the entire U.S high voltage transmission network, a better ability to diagnose failing or mis-operating equipment, and has helped prevent outages [14]. Increasing availability and familiarity with PMUs they will increase use of their data to inform a broader range of operating decisions, provided new software is developed to take advantage of the richness of synchrophasor data.

1.2 Project Overview

Bonneville Power Administration (BPA) partnered with the Rochester Institute of Technology (RIT), the University of Illinois at Urbana-Champaign (UIUC), and during the project period also with the Texas A &

M University (TAMU) to better leverage synchrophasor data to help their operators. Specifically, the scope of work was to develop new visualization solutions for PMU data to improve operator SA and decision making. The creation of visualizations was a multifaceted process which included creating a model of BPA's power system, creating scenarios to replicate situations an operator would face, running simulations of the scenarios and recording data as would be recorded by PMUs, cleaning and formatting that data, and then, finally, creating visualizations of PMU data. For clarification, we did not use data actually collected from PMUs, but instead created our own data to replicate what could be collected by BPA, so that the method was generalizable for other systems. This allowed us to test a variety of problem scenarios without the requirement that they had actually occurred in the real grid and allowed us to focus on creating visualizations instead of dealing with erroneous data from real PMUs.

The preliminary research described here, in this Part I of the final technical report, consisted of a review of relevant literature, cognitive work analysis (CWA) of the dispatchers' jobs, and simulating various scenarios for data for development of our subsequent visualization solutions. Subsequent parts (II, III, and IV) describe the results of our work.

2 Review of Relevant Literature

2.1 Current Power System Technologies

Real Time Dynamics Monitoring System (RTDMS) is developed measure the data from the PMUs. It is a software that is used for reliability measurement. Before the development of RTDMS systems, SCADA Systems were used for local monitoring and control. However, the limitation with the SCADA systems is that they are unable to capture the dynamic characteristics of the power systems, whereas the RTDMS systems are capable to deal with the high inflow volume of time-synchronized data that are delivered by phasor technology [15].

Data from the PMUs at the substations is communicated to the local Phasor Data Concentrator (PDC) which further communicates the data to the central PDC. The RTDMS integrates with the super PDC and receives all the data in a standard format, like the BPA PDCStream or the IEEE C37.118 format.

The visualization and event alarms client is in support of the real-time. The event analyzer client is used for offline support, to use the data and analyze and event in depth later to study the reasons for any accident or mishaps. The RTDMS long term analysis and report client is used for long-term monitoring and analysis functions.

To tackle the challenges of visualization of PMU data, [15] suggested that RTDMS uses the tiered visualization architecture which supports the drill-down capacity. This would be very helpful for wide-area viewing. The RTDMS also has a dashboard display that would compliment the view at the highest tier so that the operator can get a overall view of the whole system quickly. This dashboard will display all the key information in a neat and compact manner. The dashboard would indicate things like the traffic buildup in the whole system, the status of alarms (e.g. is any parameter approaching close to the alarming limit) which would be represented by different colors. Other important metrics that would be displayed in the system would include interconnection frequency, frequency instability, angle difference between flow gates and between different regions, MVAR and MW flows, voltage sensitivities, angular sensitivities, etc.

The dashboard is used to have an overview of the overall health of the complete system and keep a track of all the important metrics over wide-area viewing. If the operator detects any problem then they can use the

drill-down feature and view the concerned region or metric.

There are different graphs that can be used to represent various metrics for example, a waterfall plot can be used for a joint time-frequency domain plot which has the power spectrum magnitude on the (z-axis). Voltage stability monitoring is done using the P-V curves and Q-V curves. Voltage sensitivity scatter plot and bar chart showing the most recent voltage sensitivity trends are amongst the few charts that are used to monitor the voltage sensitivities. Other examples of chart display that is used is the Event analyzer spectral analysis display. For long term analysis and reporting the RTDMS system also generates daily performance reports and disturbance analysis reports [15].

There are several data sources that are required to get an accurate idea about the current state of the power system [16]. A few examples of the various data sources that are used are SCADA data sources, PMU data sources, transaction data sources, and weather data sources. The SCADA data source includes data such as the node voltage magnitudes, and branch flow magnitudes. The PMU provides data such as node voltage magnitudes, node voltage phase angles, standardized frequency, branch flow magnitudes, etc. Although both SCADA and PMU provide information such as the node voltage magnitudes and the branch flow magnitudes, the SCADA is a steady state system and a lot lower refresh rate than the PMU. PMU sends in time stamped data to the Phasor Data Concentrators (PDC) at the rate of 30 data measurement points per second [16].

The PMUs send the data to the PDC which stores the data and exports the measurements in real time to other various applications. Power system industry dealing with PMU measurements has a challenge of interfacing the PDC made by different companies. This is because the PDCs communicate data with one another so it is important that they should interface properly not only with each other but also with the external applications, irrespective of the vendor that built the PDC. For this purpose a technology called as the OLE for Process Control is used [16].

The idea of zooming and layering of information has been used by other systems (such as the geographic information system) but it is relatively new to the power control system and the primary reason that has been attributed to this is the reluctance to embrace new technology. One prominent reason why there is reluctance in using the technology of zooming and layering information (geographic display) is that there are several operators that have been working on the old power systems from a long time and changing it would make it difficult for them to adjust to the new changes and may affect the operators ability to do his job. The paper further highlights the various advantages of the using the geographic displays. The choice of display is very critical and could mean the difference between a good and a bad visualization [16].

A few of the interesting displays suggested by [16] are phase angle contour with superimposed voltage magnitude, the angle contours and line reclosure display helps the operator in identifying that when it is safe for line reclosure, strip charts superimposed over geographic display, rapid change displays which indicate the part or area of the system that has suffered sudden changes significantly, and worst-case overload displays indicating color coded overload in the system.

2.1.1 Situation Awareness

The August 14, 2003, blackout in Northeast US/Canada left 50 million people without power and had an estimated cost between \$4 and \$10 *billion*. The subsequent investigation cited “inadequate situation awareness” in control room for grid operations: “Training deficiencies, ineffective communications, and inadequate reliability tools and backup capabilities all contributed to a lack of situation awareness (SA) for the operators

involved.” [17]

In aviation, as much as 88% of human errors are due to problems with SA. Pilots do not get information that is needed (78%), do not correctly understand information they do get (17%), and do not project what will happen in the future (5%).

The first and still most common and generally accepted definition of SA is by Endsley [18]: “Situation awareness is the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future.” Note the three distinct *levels* of SA in this definition. Perception (Level 1), comprehension (Level 2) and projection (Level 3) are not necessarily linear stages. Levels 2 and 3 can be used to drive the search for Level 1 information. Also, default values from the mental model can provide reasonable values, even when no Level 1 info has been perceived on an element.

There is a distinction between bottom-up and top-down processing. Bottom-up, or *data driven processing* depends on external stimuli or cues in the environment (e.g., displays). Salient cues “catch” attention, are interpreted, options what to do about the cues are generated and evaluated, an option is selected, and appropriate actions are taken. In contrast, top-down or *goal driven processing* is characterized by goals that direct attention, determine development of Level 2 SA, and determine selection of a model for interpreting information.

Goals are key organizing feature for cognitive work. SA ties goals to mental models and explains dynamic goal switching and dynamic prioritization of Information. SA also facilitates selection of mental models and development of more refined schema and models through homomorphisms.

The role of SA in human performance is fundamental. According to this view, SA is fundamental to decision making and subsequent response selection (actions). On the other hand, a particular critique of SA is that it is just a “buzzword of the 90s” (and even today), a name for something we know little about, or a “black box” model of phenomena that are not understood. Nevertheless, SA is by now a “mainstream” construct and as such a useful theoretical framework to examine dispatchers’ visual information needs.

Situation awareness is important especially during the times of electrical disturbances. The Blackout of power grid takes place in two phases, (1) pre-cascading phase where the situation evolves slowly and the operators have enough time to react to the situation and (2) cascading phase where the components trip very rapidly and the operators do not have the time to react to the situation. In this case, the automatic devices come into play which protect components or try to salvage them by carrying under frequency load shedding. It is important that the operators have proper SA for them to take proper decision in the pre-cascading phase itself and prevent the stage from reaching the cascading phase. Due to the interconnectivity of the electric grid, it is becoming increasingly important that the operators also have an idea about the neighboring electric grids and control areas as well [19].

The main challenge to operators’ SA in power system control centers arises due to the dynamic nature of interaction between the human operators and between the human operators and human machine interface. The blackout of 2003 made people realize the importance of the energy management software. It is important that there is up-to-date hardware and software available so that they are able to detect the early changes in the system so that more effective response can be taken by the operators [19].

During the times of emergency and stress, the operators have to deal with a large of amount of informa-

tion and make decisions. Thus, it is important that there should be a good visualization system that has good GUI that is intuitive in nature. Intuitive GUI makes the comprehension of information a little easier and when there is a lot of information to look at a time then the operators should not miss out of the critical information needed.

Too much automation can be a problem as well, because if majority of the systems are automated then it leaves the operator out of the loop. So, in times of emergency the operator may not clear understanding of what was going on in the system when a problem occurred; we want to avoid such situation.

The operators' cognitive model acts as a filter between the *presented* power system data and the *perceived* power system state. This filter is caused due to the operators' cognitive limitations. Based on the perceived system state, the operators make changes and decide upon the future behavior of the system.

There are several recommendations for supporting SA particularly in power systems [19]:

- Accuracy of state estimation should be improved by detecting and eliminating bad analog measurement and erroneous breaker/switch status;
- In the design of the graphical user interface (GUI), the important question that should be kept in mind is that what information is really needed by the operator and when is that information needed. How should the information be organized? What is the best way to present the information to maximize the operator's performance?
- The designed system should be such that there is a good scope of human-system interaction so that the operator is in control and in the decision making loop;
- To reduce operators' individual errors there should be proper training for all operators in various scenarios so that they are adequately prepared for some emergency situation that arises;
- Wide-area monitoring and collaborative visualization should be such that the operators are able to have visualization of wide-area view boundary;
- Functionality of hardware and software applications should be ensured in all conditions and all situations.

The interface design problem may be stated as follows:

- Question: Given a complex work domain, how to describe domain complexity?
- Required: A domain representation formalism (the Abstraction-Decomposition Space, Contextual Activity Matrix, and Decision Ladders).
- Question: How to communicate the information (domain representation) to the operator?
- Required: A model of the mechanisms that people have for dealing with complexity (the SRK framework).

An answer to this problem is given by cognitive task- and work analysis, described below.

2.2 Cognitive Task Analysis (CTA)

Task analysis generally relies on observation of people performing tasks of interest. This only works when the performance of the task is *observable*, however. What about analysis of *unobservable* cognitive tasks? Cognitive task analysis (CTA) represents collection of techniques for such situations. CTA provides for description of the cognitive skills needed to perform a task proficiently. CTA, too, includes two parts, (1) knowledge *elicitation* and knowledge *representation*. CTA is valuable for tasks that depend on cognitive aspects of expertise, for example, decision making and problem solving.

The key attributes of CTA are *how* to look (interview, or self-reports, or observation, or automated capture) and *where* to look: Where in time (past, present, or future), where in realism (real-world or simulated scenarios), where in difficulty (routine or challenging tasks), and where in generality (abstract knowledge or specific events).

2.3 Cognitive Work Analysis (CWA)

Cognitive work analysis (CWA) is a framework for systematic and detailed analysis of cognitive work. As the name implies, its scope is broader than CTA, and one can see CTA as a component of CWA. CWA consists of 6 distinct stages [20].

2.3.1 Stage 1

Stage 1 is the Functional Work Structure or Work Domain Analysis (WDA), which is a technique within CWA that creates a representation of a socio-technical systems work domain, known as the abstraction-decomposition space (ADS). The ADS identifies the important, activity-independent structure of the work domain, to aid researchers in understanding the necessary values and priorities, work functions, technical functions, and physical resources to fulfill the domain purpose of the complex socio-technical system.

The purpose of an ADS is to identify aspects of a work domain that either support the achievement of the domain purpose or constrain against it. The typical ADS representation portrays the domain purpose as the final element composed of more detailed levels that follow in a hierarchical fashion.

The domain purpose is listed at the top of the representation, followed by the domain values and priorities, the work function to obtain the values and priorities, the technical functions necessary to fulfill the work functions, and ending in the physical resources required to fulfill the technical functions (either people for socio-technical systems, or technological components for technical systems). Within each of the aforementioned levels of the ADS, functions of the work domain are placed as nodes. Links between nodes at different levels represent means-ends relationships between the linked nodes.

2.3.2 Stage 2

The next stage (2), is Partitioning and Organization of Work or Work Organization Analysis (WOA). This stage focuses on domain functions, as identified in the ADS, work situations, which are the various situational contexts in which work takes place, and work tasks, which are the distinctive outcomes to be achieved. The product of this stage of analysis is a Contextual Activity Matrix.

2.3.3 Stage 3

Stage 3 is the Cognitive Transformations Analysis. It examines cognitive states established during task execution, and cognitive processes used to effect the transitions between states. The product of this stage of analysis is a suite of decision ladders, originally developed by Rasmussen [21]. Decision ladders are an extremely powerful tool to investigate just how operators perform their tasks and what information they need to do so (Fig. 1). How high on the ladder the decision maker needs to climb before arriving at an executable solution determines whether the performance is skill-, rule-, or knowledge-based. This has important implications on the design of the interface to aid in the decision-making process, as will be detailed below.

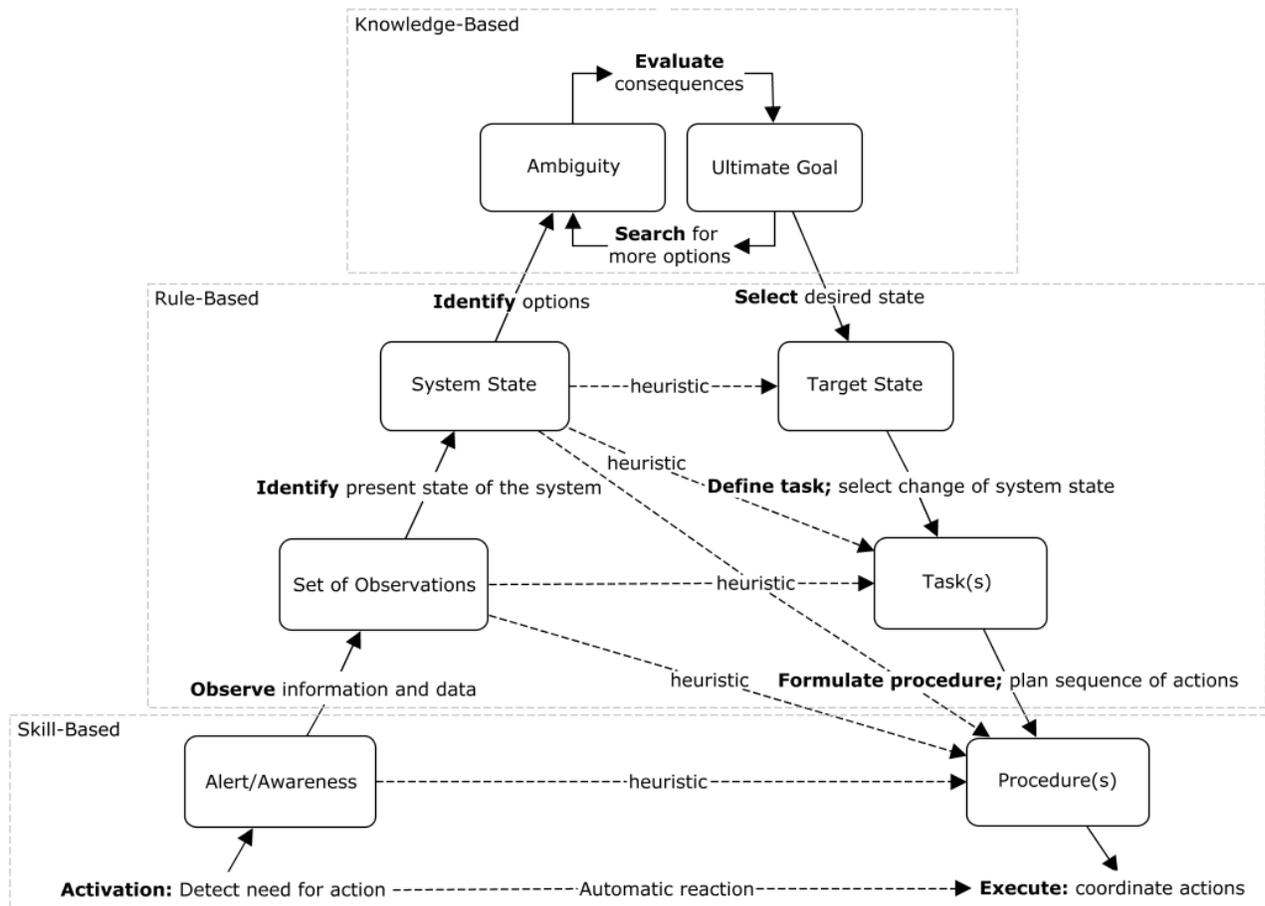


Figure 1. A decision ladder. One climbs the ladder up on the left side and down on the right side, but one may also take shortcuts afforded by heuristics without traversing the entire ladder.

2.3.4 Stages 4 and 5

Stages 4, Cognitive Strategies Analysis and 5, Cognitive Processing Analysis focus on the reasons that a worker may select one strategy in preference to another or may transition between strategies during execution of a cognitive process and identifies the skills-, rules, or knowledge-based modes of cognition [22]. In ecological interface design, the designer may choose to encourage and induce one cognitive mode over

another as dictated by the situation (e.g., in an emergency response, skill-based performance is preferred). The products of these stages of analysis are detailed description of potential strategies and of the factors that will prompt selection of one strategy over another, as well as of the activity elements associated with the different modes of cognitive processing.

2.3.5 Stage 6

Finally, stage 6, Social Transactions Analysis results in a Social Transactions Matrix, which maps agents (either human or technological or some combination) to Work Tasks and maps Work Tasks to Transaction Demands and Transaction Modes. A second product is a Transaction Network in which the transactions between agents (either human or technological) are identified and characterized in terms of fundamental or generic properties relevant to design. Figure 2.3.5 shows the 6 stages of CWA, what products follow from each process, and how each product may be used to inform design.

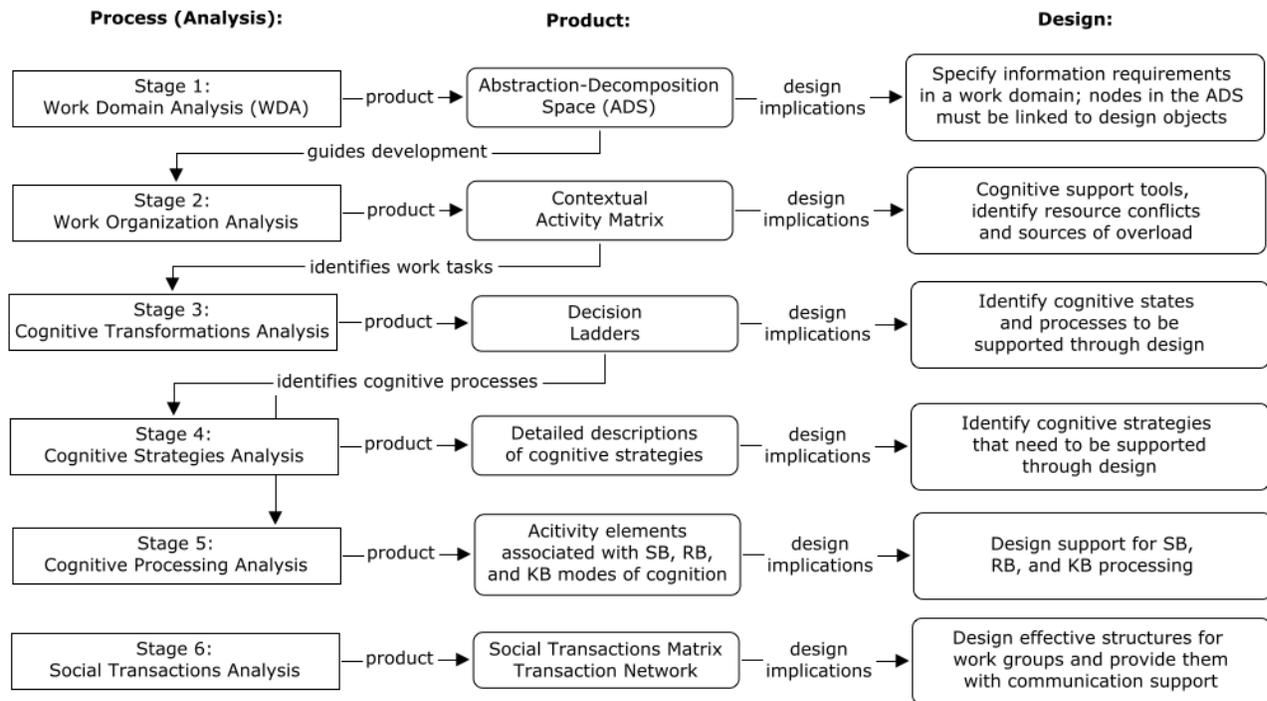


Figure 2. The 6 stages of CWA, showing the products that follow from each process, and how each product may be used to inform design.

2.4 Ecological Interface Design

The entire process of CWA is integral to *ecological interface design* (EID) [23] [24]. EID seamlessly integrates analysis, design, and evaluation functions facilitating a truly user-centered design process. EID was introduced specifically for complex sociotechnical, real-time, and dynamic systems.

The goal of EID is to make constraints and complex relationships in the work environment perceptually evident (e.g. visible, audible) to the user. This allows more of users' cognitive resources to be devoted to higher cognitive processes such as problem solving and decision making. EID is based on two key concepts

from cognitive engineering research: the Abstraction Hierarchy (AH) and the Skills, Rules, Knowledge (SRK) framework.

The design *principles* of EID fall into the SRK categories. For skill-based behavior (SBB) the operator should be able to act directly on the display and thus the structure of the display should support SBB. Direct manipulation device is therefore preferred over command-language interface. Aggregation of elementary movements into more complex routines corresponds to concurrent chunking of visual features into higher-level cues for these routines and allows for representation of multiple levels simultaneously.

For rule-based behavior (RBB) the designers should provide consistent one-to-one mapping between the work domain constraints and the cues or signs provided by the interface (i.e., provide the operators signs they can use to select appropriate actions). This will allow operators operate by relying in perceptual cues instead of having to resort to knowledge-based behavior (KBB) and to take advantage of the economy of RBB while preserving the wide applicability of KBB.

Finally, designers should support KBB by revealing the problem space in the form of an AH presentation, providing the operators a normative model of the work domain, supporting experimentation (what if scenarios) and relieving the operators of the burden of keeping track of causal nets within which they are reasoning.

2.5 Human Factors Guidelines For Display Design

Visualization of complex data in electric power systems, such as from phasor measurement units (PMUs), requires representation of numerical values of several critical variables in graphical form. Graphs concisely communicate information about relationships between between different things (variables, system components, etc.). However, to have graphical representations of data convey useful—and readily usable—*information* to human operators, who need to make timely decisions based on the displayed information, requires careful attention to the *design* of the displays. Fortunately, much good, true and tried guidelines exist that help in the design of visual displays for different kinds of data to maximize human performance in their use. This document presents these guidelines in a checklist format to facilitate their easy application to the design of PMU displays.

2.5.1 Toolbox: Graphical Variables

There is a relatively limited number of variables available for graph design, and not all of them may be contained within anyone statistical software package. Nonetheless, here is a list of the most common ones:

- **Line width:** Width of lines may be used to indicate the relative importance of graph elements. Hence, trend lines, for example, should be wider than lines depicting axes or boundaries to make them the most salient feature in a graph.
- **Line type:** Use different line types to visually differentiate between different elements (e.g., variables, groups, &c.).
- **Color (hue, saturation):** Color is a great way to attract attention to the desired elements in graphs. There are three caveats for use of color, however: (1) bright, saturated colors may add to *visual noise* and detract from the main message of the graph, (2) many publications do not allow for color graphs or may charge extra for their printing, and (3) when using color one must consider potential deficiencies of color vision among readers.

- Contrast: High contrast can be used to increase salience of key elements of the graph. Lower contrast may be used to “put” secondary information in the background where it is available to the reader but where it does not distract from the main message of the graph.
- Shape: Use different plotting symbols to visually differentiate between different groups or conditions in graphs.
- Transparency: Similarly with contrast, transparency may be used to “fade” less important information as well as allow graph elements be visible without them obscuring each other
- Text and numbers (font and size): In addition to making textual and numerical elements of the graph visible and easy to read, font and size may also be used to differentiate graph elements (e.g., larger font for primary information, smaller for secondary information).
- Animation: This option is quite advanced and not applicable to static graphs created with many statistical software applications, but it may be used in *interactive* graphs.

Although the above list is short, these graphical variables and their *combinations* allow for much creativity in display design.

The next sections provides some *general* guidelines for display design. These are not hard and fast rules but guidelines; exactly how a graph is created depends on the underlying data and what information about them is to be conveyed.

2.5.2 Hierarchy of Guidelines

The design of graphs and application of the should heed the following, simple, three-step hierarchy, with decreasing *salience* in each step:

1. The first and the most important thing is the primary message to be conveyed by the display. This should be portrayed in such a way that the reader immediately and intuitively grasps it. Use whatever graphical variables you have at your disposal to make this message salient: graph type, color, contrast, line width, &c.
2. Graph *standards* come next in this hierarchy. These include the “mandatory” elements (e.g., by APA; see below) as well as the many other guidelines detailed below. Adhering to these conventions makes use of readers’ expectations and intuition of how a graph should be read. Sometimes it may be necessary to break convention in a novel—and better—way to represent data, but a careful cost-benefit analysis is warranted in such a case.
3. Other elements. This category concerns the design of elements in graphs that are supporting rather than directly essential for the main message and that are not required by standards such as the APA Publication Manual. Examples of these elements are tick marks and numeric scales on axes, grid lines, and possible highlights (e.g., an arrow pointing to an inflection point in a line graph) to draw attention to critical pieces of information on a graph.

2.5.3 Gestalt Principles

Gestalt is a German word meaning “shape” or “form”. In perception, gestalt refers to our visual system’s ability to organize visual elements in such a way that they represent whole figures and patterns. There are several gestalt laws of grouping (or laws of *prägnanz* = “good figure”) that act as a constraint in graph

comprehension. For example, law of proximity says that objects that are close together will be perceived as a group, and the law of similarity holds that objects that share similar characteristics, such as shape or color, will be perceived as belonging to the same group.

2.5.4 3-D vs. 2-D

Never, *ever*, use 3-dimensional elements to represent 2-dimensional data. *Only* if the data presented are in fact three-dimensional, that is, must be plotted along three axes, a 3-D graph is justified (e.g., surface plots).

2.5.5 Graph Standards

There are many standards for graphs in publications. The American Psychological Association standards [25] offer very handy guidelines for creation of good graphs. By the APA standards for figures, a good figure:

- Conveys only essential facts.
- Omits visually distracting details.
- All elements are large enough with sufficient contrast to be read with ease.
- Is easy to understand; the purpose is readily apparent.
- Lines are smooth and sharp.
- Typeface is simple (sans serif) and legible
- Units of measure are provided.
- Axes are clearly labeled.
- All elements within figure are labeled or explained.

2.5.6 Other Standards

- Other conventions include depicting the *independent* variable on the x-axis and the *dependent* variable on the y-axis. It is important to heed such conventions and standards and guidelines as they represent the *expectations* of experienced graph readers. In other words, surprising the reader will result in increased risk of misunderstanding and slower graph reading.

Note that these are *general* standards. It is critical to research the conventions specific to power systems operation to create displays that are consistent with them.

- The graph should also accurately reveal the relationships as they truly are and not obscure any phenomena.
- Exploit “natural mappings” in graph design, e.g., lines for trend information, side-by-side boxes to compare two conditions or groups, or clusters of data points. Such natural mappings correspond to readers’ existing mental models and expectations and make graphs intuitively understandable. Again, there may be natural mappings unique to operators of power systems.

2.5.7 Principles of Display Design

By definition graphs are *information displays*. There are several guidelines for display design in the human factors literature that are readily applicable to graph design as well. Wickens [26] developed 13 principles of display design based on the known human capabilities and limitations. Below are 10 out of the 13 principles most applicable to graph design.

1. Make displays legible. A display's legibility is critical and necessary for designing a usable display. If the characters or objects being displayed cannot be discernible, then the reader cannot effectively make use of them.
2. Top-down processing. Information is likely perceived and interpreted in accordance with what is expected based on a user's experience. If information is presented contrary to the user's expectation, more physical evidence of that information may need to be presented to assure that it is understood correctly.
3. Redundancy gain. If information is presented more than once, it is more likely that it will be understood correctly. This can be done by presenting information in alternative physical forms (e.g. color and shape), as redundancy does not imply repetition. A traffic light is a good example of redundancy, as color and position are redundant.
4. Similarity causes confusion: Use discriminable elements. Elements that appear to be similar will likely be confused. High ratio of similar features to different features causes elements to be similar [27]. Unnecessary similar features should be removed and dissimilar features should be highlighted.
5. Principle of pictorial realism. A display should look like the variable that it represents (e.g. high temperature on a thermometer shown as a higher vertical level). We read English from left to right, and numbers on our rulers grow from left to right; low numbers should therefore be to the left of the scale and large numbers to the right there are multiple elements, they can be configured in a manner that looks like it would in the represented environment [28].
6. Principle of the moving part. Moving elements should move in a pattern and direction compatible with the user's mental model of how it actually moves in the system. This principle is applicable only to *interactive* graphs. For example, the user should move a slider *up* or to the *right* to increase the value of a variable.
7. Minimizing information access cost. When the user's attention is diverted from one location to another to access necessary information, there is an associated cost in time or effort. A graph design should minimize this cost by allowing for frequently accessed sources to be located at the nearest possible position. Displays should be so small that little scanning is required to access all information [29]. However, adequate legibility should not be sacrificed to reduce this cost.
8. Proximity compatibility principle. Divided attention between two information sources may be necessary for the completion of one task. These sources must be mentally integrated and are defined to have close mental proximity. Information access costs should be low, which can be achieved in many ways (e.g. proximity, linkage by common colors, patterns, shapes, etc.). However, close display proximity can be harmful by causing too much clutter [30].
9. Replace memory with visual information, that is, knowledge in the world. A user should not need to retain important information solely in working memory or retrieve it from long-term memory. Display of all relevant elements can aid the user by easing the use of their memory. However, the

use of memory may sometimes benefit the user by eliminating the need to reference some type of knowledge in the world. The use of knowledge in a user's head and knowledge in the world must be balanced for an effective design. Too much visual information can lead to clutter [31]. Predictive displays remove a resource-demanding cognitive task and replaces it with a simpler perceptual one [26].

10. Principle of consistency. Old habits from other displays will easily transfer to support processing of new displays if they are designed consistently. A user's long-term memory will trigger actions that are expected to be appropriate. A design must accept this fact and utilize consistency among different displays. If display layout is inconsistent, adequate alerts must be provided. Using similar color coding across several displays is helpful, so that, for example, red will mean the same thing across several displays. Display panels should be consistently organized, this reduces information access cost each time a new set is encountered [26].

3 Dispatcher Cognitive Work Analysis (CWA)

To keep human operators "in the loop" and to empower them to access relevant information and make appropriate decisions and take correct and timely actions in response to events requires that the information is displayed to them in a form that is inherently usable. Cognitive Work Analysis (CWA) is a conceptual framework that allows for analysis of all factors that affect human-information interaction. The products of this system of analyses can then be directly transformed to design requirements for information systems. The CWA approach is work-centered, rather than user-centered. It analyzes the constraints and goals that shape information behavior in the workplace regardless of the individuals who are involved in the work. It is an inherently holistic approach that simultaneously examines the environmental, organizational, social, activities, and individual dimensions. CWA provides concepts and templates for comprehensive analysis of complex phenomena but does not reduce their complexity. CWA has proven to be an effective approach to the study of human-information behavior for the purpose of designing information systems [32].

3.1 Observations

PIs Overbye and Rantanen visited the Dittmer Control Center (DCC) on December 10–11, 2015, for initial orientation and observation of the dispatchers' work domain. PI Rantanen visited the DCC again on January 26–27, 2016, and on January 30–31, 2017. Great insights were gained during two full days of observation and conversation with the DCC Dispatchers, including preliminary design criteria for PMU data visualizations.

The control room at DCC has a total of 8 workstations, or desks, of which 5 are continually manned. From left to right (facing the wall displays in front of the desks) these are the Automatic Generation Control (AGC) desk, the Remedial Actions Scheme (RAS) desk, West and East Transmission desks, and the supervisor desk in the center of the room and slightly elevated behind the aforementioned desks. The analyses relevant to the project concern only the AGC, RAS, and transmission positions and how PMU data may be used by the dispatchers manning them.

All the dispatchers emphasized the teamwork nature of the operations in the control center: The desk operators must work well together; physical proximity is important with direct communication also through gestures and expressions. The dispatchers routinely spoke to each other across the control room, highlighting the need of direct, line-of-sight, communications. The dispatchers work in 12-hour shifts without breaks. They eat and rest during periods of little activity.

The dispatchers have vast amounts of “tacit knowledge” that is not necessarily formalized in procedures. This knowledge stems from years of experience in electric power industry other than control room operation, including but not limited to substation operation, generation (hydro, coal-fired, nuclear), and line work. Many dispatchers at DCC had taken advantage of the BPA apprenticeship program, which consists of 4 years of training, from substations to line work to generation and electrical theory. Some dispatchers also had military background (e.g., in the nuclear Navy). An example of experience and tacit knowledge is a dispatcher who was observed to get a call about high voltage; he knew about it ahead of time, and as there was nothing to do about it, the call could be dealt with very quickly. Much of the work is also quite predictable. There are scheduled flows by demand by the time of day. The dispatchers know all about this and can anticipate necessary control actions. For example, the night shift gets ready for morning load increase.

We draw the following conclusions from our observations:

1. There are presently (too) many platforms, tools and displays at the DCC. Lack of uniformity or standard between myriad displays makes display upgrades difficult. A need to consolidate displays and tools is pressing.
2. Oscillation detection is one of the most important needs of dispatchers. Essentially, they only need to know *where* the oscillation is happening, or originating (e.g., at a generating station); the subsequent actions can be straightforward: call the generating station to tell them that they have oscillation and that they need to fix it. Reactors and capacitors are also used to control and stop oscillations.
3. Cascading events are a particular challenge. The dispatchers need right information and controls to halt the cascade before problems become too large to manage in a timely manner.
4. Oscillation and islanding detection, mode meters, angle alarms, MW flow, and frequency disturbance displays are useful but presently separate. Their combination into one display is a worthwhile problem to consider.
5. Much of the current efforts to integrate PMU data into the dispatchers’ “toolbox” involves automation. In other words, processing of PMU data will be done automatically with different algorithms determining alarm thresholds. For the dispatchers, this means just more, albeit presumably better, alarms to react to, and then take rule-based actions. As with all alarm systems, the reliability of alarms and proper calibration of trust on them remain an issue that will only be resolved in time, with years of experience with the new alarms.

3.2 Survey

An online survey was designed and administered to the BPA dispatchers at the Dittmer Control Center (DCC) in June and July 2015. The survey was prepared with Qualtrics online survey tool adopted by RIT and a link to the survey was forwarded to the dispatchers at DCC by Christopher Sanford, Director of Real Time Operations. Altogether 11 people responded. Because the response rate was quite poor, the results only suggest direction for the designs without providing authoritative answers to the many questions we had about the dispatchers’ information needs.

The respondents, 2 RAS/AGC operators, 3 senior system operators, 4 transmission dispatchers, and 1 working all positions, were experienced. The mean number of years in the in present position was nearly 11 years, ranging from 3 to 20 years. In the following, the results from the key questions are discussed. Please

note that the responses varied between the positions the respondents held (e.g., transmission dispatcher vs. RAS/AGC operator); with such a

3.2.1 Most Important Variables

The respondents were asked to rank 10 electrical variables in order of importance when monitoring the system status. The variables (with overall mean ranks in parentheses) in descending order of importance are:

1. Frequency (2.4)
2. Bus voltage (2.5)
3. Transmission line/transformer percentage of limit (5.1)
4. Generator real reserves (5.1)
5. Transmission line and transformer real power flow (5.4)
6. Transmission line and transformer reactive power flow (6.1)
7. Bus voltage phase angles (6.5)
8. Generator reactive reserves (6.6)
9. Interface loading (7.5)
10. Transformer tap positions (7.7)

In addition, the respondents reported the following variables they monitor: unplanned generator outages, generators bucking or boosting, unplanned outages of RAS or BES equipment, work planned for the day, if generator regulators all on auto, kinds of reactive devices (reactors and capacitors) available vs. being used, weather, symptoms of system instability, DC line in service, area control error, wind power on line, weather/lightning, lines/equipment out of service, and load vs. generation (ACE). The RAS/AGC operators ranked the variables in almost reverse order from the rest, and also came up with most of the other variables listed above.

3.2.2 Frequency of Monitoring

The respondents were asked how frequently they check the values of the following variables during real-time operation of the grid. The answer alternatives were multiple times per minute, every 5 minutes, every 15 minutes, every 30 minutes, every hour, and less than every hour. For calculating descriptive statistics these responses were given the following quantitative values, respectively: 0.5, 5, 15, 30, 60, and 90. Because these values are not on an interval scale, median values instead of means are reported below. The variables are arranged in order of frequency of monitoring:

1. Bus voltage magnitude (5)
2. Frequency (5)
3. Transmission line and transformer real power flow (5)
4. Other (5)

5. Transmission line and transformer reactive power flow (15)
6. Interface loading (15)
7. Generator real reserves (30)
8. Generator reactive reserves (30)
9. Transformer tap positions (60)
10. Transmission line/transformer percentage of limit (90)
11. Bus voltage phase angles (90)

The other variables were ACE, deployed reserves, wind power on line, weather/lightning, and system voltage overviews. These results are not reliable or usable for display design, however, for there was much variability between the respondents. Some transmission dispatchers monitored the variables much more frequently than others; the median monitoring frequency on one transmission dispatcher was every 5 minutes whereas for another it was less than every hour.

3.2.3 Decision Analysis

Respondents were asked to think of one of the most stressful operational situations they have encountered in their jobs and how they decided to act on it. The purpose of this question was to put the responses into a decision ladder for analysis. No such analyses could be performed on only four responses we received to this question. Nonetheless, the responses are interesting and illustrative, and they are offered below. Transmission dispatcher:

- **Situation description:** Relay operation with associated large RAS operation generation drop.
- **Cue:** SCADA alarms alerting relay operation of PCBs, voltage alarms from loss of MW flow on system, large negative ACE.
- **Observations:** Read the alarm and verify on the board that the expected PCBs are open.
- **System State:** SCADA, Board, ACE, input from other dispatchers identifying.
- **Options:** Adjust voltage, call power house and energize line back to them if they are ready; test energize line that relayed out of service
- **Evaluation of options:** We have a DSO for testing of faulted equipment. Ensure no equipment alarms on equipment I will remotely operate. Ensure unit PCBs are open at generation facility prior to energizing line back to them.
- **Select option:** Confirm with coworkers that they see the same indications I see and have no objections to my actions.
- **Target state:** Stable voltages, faulted lines back in service, generation having a hot line to synch to.
- **Tasks:** Test energizing line, energizing line back to generation facility

RAS/AGC operator:

- **Situation description:** Hot weather and over on COI and other cut planes like North of Hanford when there were fires.

- **Cue:** Alarm plus my observation on turnover of weather conditions, previous shift conditions, loading on paths.
- **Observations:** I looked at the alarm screen after alarm went off, went to path loading pages to see actual flows, how much are we over, briefly investigate if there is a sudden unplanned outage of transmission, generation or load on the system and if it is momentary or will be out long.
- **System State:** When I have the knowledge from the previous question then I determine whether I need to take action and how fast I need to take action.
- **Options:** Get equipment back in service if necessary, contact CISO to see if they are able to mitigate if it is their problem, cut schedules, circulate and coordinate on the PDCI.
- **Evaluation of options:** How much am I over, how much time do I have to act, will this be an ongoing problem in future hours, how cooperative are other entities and are they able to respond, what is the impact to the water schedule. In my mind i am going through every option and when I get to a point where it will take too long to do something, I go on to the next item and follow my thought process through.
- **Select option:** I choose the one that will have the least impact to schedules, generation, and has the least impact to other entities, the option that will definitely solve the problem; is the hammer big enough (if you are over by 100 MW, you have have to take action for 150 MW because if you only take action for 100 MW, it might not be enough and you will not accomplish the relief in the appropriate amount of time).
- **Target state:** Under path limits and not likely to continue to go over the limit.
- **Tasks:** Notifications, who do I need to notify or work with, what if they cannot or will not cooperate, what if the decision i make is not enough, what is my Plan B, am I creating another problem with my choice of solution?
- **Procedures:** Decide with senior if necessary, notify coordinating entity and agree on solution, what generation will I move, who will I notify, prioritize tasks (I cannot just make notifications, there has to be some action).
- **Actions:** Circulated power to off load the path and cut schedules for future hours.
- **General comments:** Particular difficulties were possibly missing an arming point, failovers of equipment causing problems and missed issues, and not enough awareness of consequences. When acknowledging alarms, you might inadvertently lose sight of an alarm that was legitimate and needed to be addressed. Alarm should silence but not disappear. During failovers, some quantities get “stuck”, and alarms get lost. Why can we only see what BPA’s internal load is by going to the manual load shed page? Why is that not a more prominent? It is an “at a glance” number that can give a quick assessment of one aspect of how loaded my system is. A temperature/fire overlay option for an overview page [would be good]. Allow me to see color (not just black) in areas where I am having extreme temperatures or problem fires or storms moving in so I can see my voltages, but I can also see what other mother nature problems are moving through. I would like to see a system wide page that was on a map and showing RAS areas so different RAS schemes that cover different areas would overlay the system map and show which schemes overlap. Make it possible to change what kind of alarm SCADA produces by tone, pitch, etc. Also, instead of a black background, when we are looking at voltage overview pages, show them overlaying a map of the northwest.

RAS/AGC operator:

- **Situation description:** Exceeded path limitations with California.
- **Cue:** System alarm.
- **Observations:** Verified over limits.
- **System State:** Compared limit with actual flows.
- **Options:** Curtailments

Note that here the operator took the shortcut to desired target state without evaluating and selecting options

- **Target state:** Below path limit.
- **Tasks:** Use curtailment tool.
- **Procedures:** Implement curtailment.
- **Actions:** Notify BC Canada to back off flows on our Northern Intertie.
- **General comments:** Too many tools to choose from that do not have enough visual impact to make intuitively obvious decisions. Numbers do not have the visual impact graphics do. Some screens have so much information it is difficult to find what you are looking for quickly. A large screen that you could instantly see the flows and changes on the system [would be desirable]. Make the screens dark for ease on the eyes with graphics that immediately show you the state but not too busy to confuse.

Senior System Operator:

- **Situation description:** SOL exceedance.
- **Cue:** Alarms, [both] audible, visual.
- **Observations:** Are there any other alarms/events, magnitude of exceedance.
- **System State:** Alarm summaries, map board, SCADA.
- **Options:** Do nothing, contact others for action, take action to mitigate.
- **Evaluation of options:** Causal-based.
- **Select option:** Followed standing order guidelines.
- **Target state:** Operations below the SOL.
- **Tasks:** Moving generation/cutting schedules.
- **Procedures:** Who to call, tools to use, timeframe available.
- **Actions:** Called for generation redispatch, reduced schedules.

- **General comments:** It is difficult to get a good system overview. There are many alarms during events. It is reactionary no predictive. It does not fully give visibility for control center backup functionality. The system was experiencing oscillations. The SCADA system did not provide granular enough information to detect the oscillations. Displays that bring together relevant, usually disperse, information for quick and accurate assessment [would be desirable], i.e., voltages, flows, limits, times, maybe resources available to mitigate/assess. There has been resistance in the past to implementing new software packages. This resistance is understandable in that some organization, which is already very busy, would have upkeeping of this software added to their duties.

3.3 Summary of CWA

The key to a successful CWA is engagement with the subject matter experts (SMEs) in the work domain of interest. We had several very valuable opportunities to observe the dispatchers at the Dittmer control center (DCC) and speak with them about their work, different tasks they do, and the tools at their disposal. These observations provided a good foundation for the CWA. However, subsequent attempts to “fill in the blanks” of the CWA were hampered by our inability to collect more detailed data from the DCC dispatchers. Only 11 dispatchers responded to an online survey and their responses were too varied for clear conclusions. Nevertheless, the CWA affords the following summary as a basis for our subsequent efforts in PMU data visualization:

- The dispatchers primary mode of operation is rule-based. There are two reasons for this: (1) The work is necessarily proceduralized to an extreme degree, and procedures necessarily follow the IF–THEN formula, which is also a hallmark of rule-based behavior. (2) The work is largely alarm-driven, or reactive. Alarms, which are a necessity in operation of a complex and highly dynamic system such as the electric grid, are tightly coupled with the actions they should trigger. This again makes alarm-driven work largely rule- or even skill-based. The goal of CWA is to determine to what degree work is skill-, rule-, or knowledge-based; despite the incompleteness of the CWA we were able to perform, we are satisfied to reach this goal nevertheless.
- The dispatchers have also largely internalized the rule-based mode of operation. There were several newly developed synchrophasor displays [33] at the DCC when we visited, but the dispatchers did not use them. In fact, the dispatchers were quite apprehensive of giving the video-wall real estate to what they called “useless” synchrophasor displays . When interviewed further, the dispatchers revealed that this verdict arose from the lack of procedures linked to the the information displayed. In other words, the dispatchers did not know what to do with the information they gleaned from the synchrophasor displays because they had not been told what to do with it.

4 Design Principles

As we saw with the synchrophasor displays at DCC, it is very difficult to introduce new tools in operational environments. Operators are rightly wary of making critical decisions based on visualizations they may not fully understand and the reliability of which is still untested. However, the visualization solutions developed within this project propose an even more radical shift in the operational culture at DCC, away from rule-based, or reactive, mode to knowledge-based, *proactive* mode. Based on the CWA, as far as it could be completed, the following principles were heeded in the design of PMU displays:

1. Given the importance of teamwork and shared SA at the DCC, the primary use of PMU data is to provide dispatchers an overview of the overall system status, displayed on the video-wall for easy (i.e., minimized information access cost), simultaneous, viewing by all.

2. Given the complaints about reactive, alarm-driven, mode of performance presently dominant at DCC. The “grid events” that are manifested in the PMU data are very rare and very sudden and very transient. As such they are difficult to detect and therefore use of automated alerts is justified. On the other hand, extensive automation has many problems with misses and especially false alerts, in terms of the signal detection theory. The dispatchers also complain about “alarm fatigue” in their jobs.
3. New PMU data displays should allow for continuous monitoring of the overall system status and *early* detection (i.e., preceding alarms) of any disturbances that might indicate problems later in time. Such displays would allow dispatchers shift from reactive to *proactive* mode of performance, or from skill- and rule-based to knowledge-based performance, and tactical and strategic control of the system. Novel displays could also allow for detection of *precursors* to events that are too complex or too small for automated alerting. Such shifts would also distribute workload more uniformly and result in overall reduction of workload.
4. The key problem to be solved is depiction of time on a visual display. There are only 4 ways of doing this: (1) devote an axis on a data plot to time, that is, plot the variable of interest against time, (2) animate the display so that the movement of the display elements in time convey the temporal information, and (3) plot derivative variables that in themselves embed time, such as rate of change (but of course, rate of change also varies in time) or see if the Fourier transformation of the data may represent time in a meaningful way.
5. A corollary problem with plotting variables of interest against time is scaling of the time axis so that the time depicted on a display matches the dispatchers’ monitoring behavior (i.e., for how long should a transient event persist on a display so that the probability of a dispatcher seeing it acceptable?). There are multiple solutions to explore, including a logarithmic scale for the time axis so the much historical information is preserved for as long as possible but current and near-past events are displayed in greater detail.
6. Scaling of time information in animated displays is a somewhat more complex problem. Potential solutions to be explored include “looping” of the animation so that the loop advances some increment (i.e., loop length) each time, as well as the total length of the loop and speed of the animation (real-time, slow-motion, or sped up).
7. Plotting of derivative variables such as rate of change and Fourier-transformed data were investigated.
8. Displaying multiple, related, variables on a single display or multiple displays arranged in close proximity to each other to take advantage of emergent features were investigated.

5 Simulated Scenarios

Several High-Impact, Low-Frequency (HILF) grid events, or a class of risks that have potential to cause catastrophic widespread impacts on the electric power system, but either rarely occur, or, in some cases, have never occurred, and thus operators have little experience dealing with them, were created to generate PMU data for development of different visualizations. The scenarios are comprised of initiating events and subsequent system events. Offering better visualizations for these events is an important selling point for power system visualization research.

Examples of HILF events include coordinated cyber, physical, and blended attacks, high altitude detonation of a nuclear weapon (EMP), natural disasters (earthquake, tsunami, large hurricanes), and geomagnetic disturbance (GMD). The initiating event (initiator) is the manifestation of a hazard that results in a sequence

of system events—failures or successes—that force the grid to an undesirable state. System events are the repercussions of the initiating event that occur in the power system, which can affect buses, circuit breakers and relays, generators, loads, transformers, transmission lines, and transmission towers.

To fully represent BPA’s system in the WECC and to limit the amount of data produced during simulations, and thus computational expenditure, equivalent models of the WECC were created. These models are also referred to as “cases” and are made to work within PowerWorld. We created an equivalent power system model of the WECC that includes all buses with an area name of “Northwest”, several large border buses outside of that area, and all generators (not in “Northwest” area) greater than 300 MW. We created a second equivalent, from the equivalent as described above, that uses only buses with nominal voltages above 230 kV and generators greater than 50 MW. Table 1 lists the scenarios created with brief descriptions.

5.1 Simulated Scenario Descriptions

To create data visualizations and then test their effectiveness, power system data was needed. To generate data, situations that mimic those that an operator might face in real life were created. These scenarios were designed to replicate power system events that can arise from storms, attacks, or other physical equipment damage. For each scenario, two simulations were created that lasted 60 seconds and 120 seconds, and were run in a power system software simulation package. Frequency, voltage angle, voltage magnitude, and current magnitude were measured at each bus in the 1,300 bus equivalent model, 30 times per second. These types of data values are the same as those that could be recorded by a PMU. Current angles were not considered, partly due to time constraints and partly because we felt they would not provide extensive information beyond what was already known from voltage angles. Data was saved in separate csv files corresponding to its type—frequency, voltage magnitude, voltage angle, and current magnitude.

Within the software simulation package we used, there is a limited number of settings that can be adjusted to create problematic events for a power system model. For example, there is not a way to tell the program to induce an earthquake on a particular area of a system. So, to create a problematic event, like an earthquake, we had to choose from settings available in the package. These adjustable settings include increasing or decreasing the amount of demand or generation in a model; faulting (short circuiting) certain buses, lines, or transformers in a model; disconnecting one or more of a model’s buses, generators, loads, lines, and transformers; or re-connecting one or more of a model’s out-of-service buses, generators, loads, lines, and transformers. The combination of settings we used to create four problematic scenarios is briefly described in the following subsections, so they can be replicated in other non-BPA systems. Though, in order to protect the security of BPA’s system, the following scenario descriptions are purposefully vague so as not to reveal information that would compromise BPA’s electric infrastructure.

5.1.1 Substation Attack

Motivated by the California Metcalfe incident, several transformers were taken out of service at one of BPA’s largest substations. Nothing else was done besides taking the transformers offline.

5.1.2 Earthquake

Buses and transmission lines were faulted and/or removed from service in order to mimic what could happen during an earthquake along Oregon and Washington’s Pacific coast. Line faults were used to simulate real transmission lines being shaken to the ground or to another phased line. It was assumed that most of the faults could not be cleared because they were lying on the ground or touching another line. As would be done in the real grid for persistent faults, the faulted lines were opened to stop the short circuit. Transformers

were opened to simulate physical damage caused by shaking. In Figure 2.2, the area to the left of the black line shows where the most damage was incurred.

5.1.3 Ice Storm

Transmission line faults and openings were used to mimic what could happen during an ice storm. Wind and/or the weight of ice was assumed to weigh down transmission lines, such that they broke from their towers and fell to the ground, causing line faults. Just like in the earthquake simulation, for faults that could not be cleared, the affected transmission lines were opened.

5.1.4 Geomagnetic Disturbance

The software simulation package includes a tool to simulate geomagnetic disturbances, which was used to create a geomagnetic storm to act on the equivalent model. Additionally, several large transformers were opened during the simulation as if they had overheated from high currents induced by the geomagnetic storm. Several loads were shed too, as might occur during a real geomagnetic storm. Varying field strengths of 6 V/km at 60 degrees, 8 V/km at 75 degrees, and 10 V/km at 80 degrees were applied to create the storm. It was not paramount that the storm be totally realistic, but instead that it provide what an operator could see during a geomagnetic storm.

Table 1.

Summary of problem scenarios that were created and used in simulations of a BPA system model.

Scenario Type	Description
Earthquake	Transmission line faults and openings, transformer faults and openings, bus faults and openings
Ice storm	Transmission line faults and openings
Physical Attack	Transformer damage (openings)
PMU Measurement Error	Slightly skew voltage angle and magnitude data
PMU Cyber Attack	Inject false voltage magnitude or angle data
Geomagnetic Disturbance	Large current flows, transformer openings, and load shedding
Generator Outage	Take a large generator offline
Load Shed	Decrease demand at certain buses/substations
Open AC Transmission Lines	Open large AC transmission line
Open DC Transmission Lines	Open DC transmission lines

5.2 Data Collection and Processing

Once the equivalent model and simulations were completed, the data collected from each simulation was processed. Data processing included (1) formatting the data to have uniform and concise identifiers and (2) extracting data that corresponded to what PMUs would collect in the real system. To automate data processing, a software tool was created by the University of Illinois team. The tool was built to work for other simulations or power models beyond this project, so as to streamline the process between simulation data and analysis. The data processing functionality of the software tool is described in this section, and a complete user guide for the tool is given in the appendix.

5.2.1 Simulation and Model Data and Formatting

Measurements for frequency, voltage magnitude, voltage angle, and current magnitude were recorded at every bus in the model system, 30 times per second (30 Hz). Data for each of the aforementioned measurement type was saved, directly from the software simulation package into a csv file. Frequencies were measured in Hertz, voltage and current magnitudes in per unit, and angles in degrees. The aforementioned software tool was used to adjust formatting of each dataset, so that bus numbers were used as the column names and the timestamp for each observation as the row name.

In addition to measurement data, information about the power system being modeled in the simulation package, including information for buses, substations, and generators was saved in separate csv files. These data, like the measurement data, was re-formatted and all extraneous information discarded. The pertinent information kept for each power system component was, for buses, the bus name, number, nominal voltage level, and substation it belonged to; for generators, the bus number it was connected to, the substation it belonged to, and its MW generating capacity; and for substations, the substation identifier (number), operating area it belonged to, and geographic latitude and longitude coordinate. Substation and bus information were merged to form one dataset that contained the bus name and number, substation identifier, nominal voltage, operating area name, and geographic latitude and longitude coordinates.

The goal of this step was to make more concise variable identifiers and a uniform format between measurement types and model information. The resulting data were considered the complete dataset, from which PMU data were derived, as is described below. The appendix provides more information about the data input to and output from the software tool.

5.2.2 PMU Data Processing

As mentioned above, measurements for frequency, voltage magnitude, voltage angle, and current magnitude were recorded at every bus in the model system. However, for our research, we were only interested in data that would be provided by PMUs, which, in general, are not installed at every bus in a system. To reconcile this, an algorithm was created to reduce the number of buses from which data was collected, to more realistically represent what would be measured by PMUs. The PMU data processing algorithm is a two-step process, which begins by creating a list of substations and buses where PMUs are actually located or where they should fictitiously be assigned. The second part of the algorithm extracts data from the complete set of simulation data, for only the buses included in the list made in step 1.

In the first step of the two-part data processing algorithm, substations where PMUs were known, or were likely, to be located in the real system, were identified, and then a list of buses contained within those substations was generated. To do this, the algorithm first looked for the user-defined number of substations where PMUs were actually installed in the real power system, or the number of PMUs desired by the user. That number is represented by N . Next, the algorithm checked if there was a csv file that contained information where PMUs were actually located. The contents of that csv file were the substation name, in the first column, and the nominal substation voltage, in the second column. If there was such a file, the software tool assigned the number of known PMUs, in that file, to a value, K , and the substation and its corresponding information (identifier, voltage level, and geographic coordinates) to a data structure, `pmu_subs`. Otherwise, K was set to zero and `pmu_subs` remained empty. Substation information was gathered by cross checking the known substation identifiers, with the model data.

With K assigned, the program then checked if the number of known PMUs (K) was equal to the required

number of PMUs (N). If so, the PMU substation identification and assignment process was complete. If not, the difference in known and needed PMUs was calculated, and the result was the number of PMUs that needed to be assigned (A) by the program, such that $N = K + A$.

When A was non-zero, the assignment process for unknown PMU locations was initiated. It identified substations that were most likely to have PMUs installed and fictitiously assumed there to be a PMU at each of those substations. The approach was based on the notions of [34], which suggests that PMUs are most often found at substations with a generator and which have buses operating at the highest and most widespread system voltages. Widespread means that if there happens to be only a few buses at a high voltage, the program will ignore those, and instead use the next lesser voltage level that has a greater number of buses. In BPA's case, there were a few buses at 345 kV, but a many more at 230 kV and 500 kV, so those were the two voltage levels used for PMU assignments. We assume PMUs monitor only two voltage levels.

The PMU substation assignment approach is described by the numbered steps:

1. Identify the two highest and most widespread voltage levels in the power system model. This is done automatically by the algorithm or can be defined by the user. In BPA's case, the two highest and most widespread voltages levels are 500 kV and 230 kV.
2. Sort all substations into categories based on their voltage level and whether or not they have a generator. Once sorted, arrange the substations, with generators, in order of the largest MW generating capacity and all other substations alphabetically. Only substations with buses at the voltage levels identified in Step 1 are considered. *Note:* Alphabetical ordering for non-generator substations was used to ensure repeatability between simulations, and is the default for the program. If repeatability is not desired, no ordering need be done, but modifications to the program would need to be made.
3. Assign PMUs to each substation that has both a generator and buses operating at the highest and most widespread system voltage, e.g., 500 kV in BPA's system.
4. If all PMUs are assigned, the process is complete, and all substations with a real or fictitious PMU are contained in `pmu_subs`. If not, assign PMUs to each substation that has both a generator and buses operating at the second highest and most widespread system voltage, e.g., 230 kV in BPA's system.
5. If all PMUs are assigned, the process is complete, and all substations with a real or fictitious PMU are contained in `pmu_subs`. If not, assign PMUs to any remaining substations that contain buses at the highest and most widespread system voltage, e.g., 500 kV in BPA's system.
6. If all PMUs are assigned, the process is complete, and all substations with a real or fictitious PMU are contained in `pmu_subs`. If not, assign PMUs to any remaining substations that contain buses at the second highest and most widespread system voltage, e.g., 230 kV in BPA's system. At the conclusion of this step, all substations with a real or fictitious PMU are contained in `pmu_subs`.

At each step, the substation, where PMUs were fictitiously assigned, and its corresponding information, were saved to the data structure, `pmu_subs`. Note that a PMU will only be assigned to one unique voltage level, so if a substation has buses at both 230 and 500 kV, two PMUs would be needed. The assignment process was complete once the number of substations assigned was equal to A and consequently the number of substations in `pmu_subs` was equal to N . For BPA, there were 126 PMUs installed [35]. Of those 126 PMUs, 56 substations were known to have PMUs from the public document of [36], but the remaining 70 had to be assigned, since their exact locations were not publicly known. Note, however, that the location of substations in BPA's system is public, shown on their website.

With the dataset of PMU substations complete, a more granular list of buses that should be monitored by each PMU needed to be created. To do this, we assumed that each PMU had enough channels to monitor all buses in a substation, for a unique voltage level. Then, for each substation in `pmu_subs`, all buses within that substation that operated at either the highest and/or second highest voltage were saved to a list called `pmu_buses`. Once the list of buses was made, it was used in the next step of the PMU data processing algorithm, for data extraction.

The main goal of the PMU algorithm was to generate realistic data like that collected from PMUs in an actual system. To do this, the number of buses from which measurements were collected had to be reduced to represent only what would be collected by PMUs. In Step 1 of the PMU data processing algorithm, a list of buses that the PMUs were monitoring, called `pmu_buses`, was created. The second step is then to create the actual set of PMU data.

The `pmu_buses` list was used to select from which buses, of the complete set of simulation data, measurements should be extracted. In other words, any bus contained in both `pmu_buses` and the data set for frequency, voltage angle, voltage magnitude, or current magnitude, was saved in a new data structure to represent PMU data. Recall that simulation data included frequency, voltage angle, voltage magnitude, and current magnitude at every bus in the model.

To extract only PMU data from the simulation data, the buses (columns) that are shared between the frequency simulation data and buses in `pmu_buses` are identified. Data from those columns are extracted and saved to a new data structure. The resulting data structure is saved in a csv file with the naming convention of `pmu_datatype`, where `datatype` is one of “freq”, “vang”, “vmag”, or “cmag”. This data extraction process must be performed for each measurement type (frequency, voltage angle, voltage magnitude, and current magnitude).

The data processing and PMU identification, assignment and data extraction process works for any power system, not just BPA’s, and is contained within the custom software tool we created. It is flexible enough to accommodate cases where PMU locations are known and cases when they are not. From here on, any time the terms PMU or PMU data are used, it is with regard to the N real and/or fictitiously assigned PMUs generated in this process, unless otherwise connoted by the context. PMU data refers to the data saved in `pmu_freq`, `pmu_vmag`, and `pmu_cmag`.

References

- [1] W. A. Wulf. Great achievements and grand challenges. *The Bridge*, 30(3-4):5–10, 2000.
- [2] G. Bakke. *The Grid: The fraying wires between Americans and our energy future*. Bloomsbury Publishing USA, 2016.
- [3] DOE Grid Tech Team. Vision of the future grid. <https://energy.gov/under-secretary-science-and-energy/vision-future-grid>.
- [4] F. Li, W. Qiao, H. Sun, H. Wan, J. Wang, Y. Xia, Z. Xu, and P. Zhang. Smart transmission grid: Vision and framework. *IEEE transactions on Smart Grid*, 1(2):168–177, 2010.
- [5] U.S. Department of Energy. The smart grid: An introduction. <https://energy.gov/oe/downloads/smart-grid-introduction-0>, 2008.

- [6] Y. Aillerie, S. Kayal, J-P. Mennella, R. Samani, S. Sauty, and L. Schmitt. Smart grid cyber security. *White Paper, Intel Corporation, McAfee, and ALSTOM*, 2013.
- [7] E. Dall’Anese, P. Mancarella, and A. Monti. Unlocking flexibility: Integrated optimization and control of multienergy systems. *IEEE Power and Energy Magazine*, 15(1):43–52, 2017.
- [8] L. Bainbridge. Ironies of automation. In K. Duncan J. Rasmussen and J. Leplat, editors, *New technology and human error*. Chichester: John Wiley Sons, 1987.
- [9] M. R. Endsley and E. O. Kiris. The out-of-the-loop performance problem and level of control in automation. *Human Factors*, 37(2):381–394, 1995.
- [10] R. Parasuraman and V. Riley. Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2):230–253, 1997.
- [11] R. Parasuraman, T. B. Sheridan, and C. D. Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 30(3):286–297, 2000.
- [12] R. H. Miller and J. H. Malinowski. *Power system operation*. McGraw-Hill Book Company, New York, NY, 3rd edition, 1994.
- [13] P. W. Sauer, M. A. Pai, and J. H. Chow. *Power System Dynamics and Stability: With Synchrophasor Measurement and Power System Toolbox*. John Wiley & Sons, 2017.
- [14] U.S. Department of Energy, Office of Electricity Delivery and Energy Reliability. Advancement of synchrophasor technology. https://www.smartgrid.gov/files/20160320_Synchrophasor_Report.pdf, March 2016.
- [15] M. Parashar and J. Mo. Real time dynamics monitoring system (RTDMS): Phasor applications for the control room. In *Proceedings of the HICSS’09, 42nd Hawaii International Conference on System Sciences*, pages 1–11. IEEE, 2009.
- [16] R. Klump, R. E. Wilson, and K. E. Martin. Visualizing real-time security threats using hybrid SCADA/PMU measurement displays. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS’05)*. IEEE, 2005.
- [17] U.S.-Canada Power System Outage Task Force. *Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations*. 2004.
- [18] M. R. Endsley. Situation awareness global assessment technique (sagat). In *Proceedings of the IEEE 1988 National Aerospace and Electronics Conference (NAECON)*, pages 789–795. IEEE, 1988.
- [19] Mathaios Panteli and Daniel S Kirschen. Situation awareness in power systems: Theory, challenges and applications. *Electric Power Systems Research*, 122:140–151, 2015.
- [20] G. Lintern. The foundations and pragmatics of cognitive work analysis: A systematic approach to design of large-scale information systems. <http://www.cognitivesystemsdesign.net/home.html>, 2009.
- [21] J. Rasmussen, A. M. Pejtersen, and L. P. Goodstein. *Cognitive systems engineering*. Wiley, 1994.
- [22] J. Rasmussen. Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man and Cybernetics*, 13(3):257–266, 1983.

- [23] K. B. Bennett and J. M. Flach. *Display and interface design: Subtle science, exact art*. CRC Press, 2011.
- [24] C. M. Burns and J. Hajdukiewicz. *Ecological interface design*. CRC Press, 2013.
- [25] APA. *Publication Manual of the American Psychological Association*. American Psychological Association, 6th edition, 2010.
- [26] C. D. Wickens, S. E. Gordon, and Y. Liu. *An Introduction to Human Factors Engineering*. Longman, New York, 1998.
- [27] A. Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977.
- [28] S. N. Roscoe. Airborne displays for flight and navigation. *Human Factors*, 10(4):321–332, 1968.
- [29] P. Kroft and C. D. Wickens. Displaying multi-domain graphical database information: An evaluation of scanning, clutter, display size, and user activity. *Information Design Journal*, 11(1):44–52, 2002.
- [30] C. D. Wickens and C. M. Carswell. The proximity compatibility principle: its psychological foundation and relevance to display design. *Human Factors*, 37(3):473–494, 1995.
- [31] D.A. Norman. *The Design of Everyday Things*. Basic Books. Basic Books, 2013.
- [32] R. Fidel and A. M. Pejtersen. From information behaviour research to the design of information systems: the cognitive work analysis framework. *Information Research*, 10(1):10–1, 2004.
- [33] BPA. Synchrophasor display overview. Technical Report Version 1.0b, Bonneville Power Administration, 2015.
- [34] L. Beard, M. Patel, P. Quinn, A. Silverstein, D. Sobajic, and L. Vanfretti. Guidelines for siting phasor measurement units. *North American SynchroPhasor Initiative*, 2011.
- [35] BPA.gov. Synchrophasor success lands bpa its first platts award. <https://www.bpa.gov/news/newsroom/Pages/Synchrophasor-success-lands-BPA-its-first-Platts-Award.aspx>, 2013.
- [36] U.S. Department of Energy, Bonneville Power Administration. Environmental clearance memorandum. https://www.bpa.gov/efw/Analysis/CategoricalExclusions/cxl/CX-ProgrammaticSynchrophasorGPSantennas_WEB.pdf, 2011.

Part II

Spatial–Temporal Visualization of PMU Data

1 Clustering

1.1 DBSCAN Clustering

Clustering with the simulated PMU data was used for two reasons: (1) to reveal any patterns or features hidden in the data, and (2) as a way to condense a large amount of data into summaries of it. For operators, that big picture perspective is preferable to more detailed knowledge about every component in a system [1], which contributes to the common operational challenge of information overload.

DBSCAN clustering is a technique that sorts data into groups that are densely packed together, and those that are not. A detailed procedure and visualization of the DBSCAN algorithm is provided in [2], from which the following summary was created. The average run time complexity for the DBSCAN algorithm is $O(n * \log(n))$, where n is the number of data points being clustered [3].

Before the algorithm begins, the user must define two numbers. The first is a number that specifies how many points are required to be near a point, in order for that point to be considered a part of a cluster. This number is called `min_points`. The second number, `epsilon`, tells how far away from the original point the algorithm can look for other points to meet the `min_points` requirement. With the parameters for `min_points` and `epsilon` defined, the algorithm begins by selecting an arbitrary data point from the dataset. From that data point, the algorithm looks for points within a radius of `epsilon`. If it finds at least `min_points` number of points within that radius, including itself, all those points are added to a cluster. If there is a point that does not have at least `min_points` around it, within a radius of `epsilon`, that point is considered noise.

Then, for each of the points that were just added to the cluster, the algorithm looks for at least `min_pts` number of points, within a radius of `epsilon`, just as was done for the initial arbitrary point. This process repeats for each point iteratively added to a cluster, until there are not any more points that have at least `min_points` around them, within a radius of `epsilon`. In other words, once the cluster surrounds itself by noise points, the cluster is complete.

Upon completion of a cluster, a new arbitrary point is chosen from the data set to see if a new cluster can be formed. The process described in the previous paragraph is repeated for this new point and the subsequent ones that are grouped with it. It is possible that a point that was in a previous cluster can be assigned to a new cluster. That is not a problem. The formation of new clusters occurs until all the data has been classified.

For our application, we used DBSCAN as a way to determine the likely number of clusters that could be formed, not really for the cluster assignments it creates. Instead, K-means was used for classifying data.

1.2 K-Means Clustering

K-means clustering, also known as Lloyd’s algorithm, is an iterative technique used to classify groups of similar data. It is a widely used clustering technique that works well for large amounts of data. One challenge with K-means is that the number of clusters, or groups, must be specified before running the algorithm. The documentation of [4] addresses that problem, and we provide a solution in the next section. The K-means algorithm is summarized in the remaining paragraphs of this section, and is fully documented in [4].

First, n samples are chosen from the original data, where n is the number of clusters that should be formed. One sample is the same as one column (variable). These n samples are used as the initial centroids, around which clusters will be formed. It should be noted that the choice of initial samples can affect the number of iterations and convergence of the algorithm [4].

Second, each column (variable) of the original data set is assigned to the nearest centroid. Nearness is determined by calculating the Euclidean distance between each column and centroid. The smaller the Euclidean distance, the nearer the column of data is to a centroid, and thus will ensure its assignment to that centroid. The Euclidean distance formula is given by

$$d = \sqrt{\sum_{i=1}^n (p_i - q_i)^2} \quad (1.1)$$

where d is the Euclidean distance, p_i is the value in row i , of a column of data, q_i is the value in row i , of the centroid column of data, i is the iterator specifying the row number, and n is the total number of rows of the data. In other words, each column of data is considered to be the vector, $\mathbf{p} = (p_1, p_2, p_3, \dots, p_n)$, with the subscript identifying the row number of each data point (p) in the column. The same vector definition holds for the centroid, where the column of centroid data is the vector $\mathbf{q} = (q_1, q_2, q_3, \dots, q_n)$.

To determine which cluster a column belongs in, we need to determine the Euclidean distance between itself and each centroid. If the Euclidean distance between column and, say, centroid 1 is the smallest, the column is assigned to cluster 1. This process is carried out for each remaining column of data, until all columns have been assigned to the nearest centroid.

New centroids are calculated by taking the mean of all samples, per row, in each newly formed cluster. Finally, the difference between the old and new centroids is compared to a pre-specified tolerance value. If the difference exceeds the tolerance, steps 2, 3, and 4 are repeated until the centroid difference is within the specification, or until a maximum number of iterations is reached. If the specification is met, the clustering process is complete, and the final clustered data is ready

The average complexity required to perform K-means is given by $O(k * n * l * T)$, where k is the number of clusters being formed, n is the number of variables (columns, buses), l is the number of observations per variable (measurements at each time step), and T is the number of iterations the algorithm takes to converge [4].

1.3 Combined Clustering Approach

To consolidate data to provide a summary of the state of the electric grid, a two-staged clustering approach was used for each simulation. Simulation data that was clustered included frequency, voltage angle, and voltage magnitude. Due to time constraints, clustering was not performed on current magnitude data.

Originally, various clustering techniques were tested to see which might be the most effective in identifying problematic buses and which provided the best overview of the system. It was found that DBSCAN was useful for identifying the number of clusters that could be formed and K-means did well classifying similar buses and highlighting the problematic ones. As such, a combination of the two techniques was used to leverage their respective strengths.

The clustering process was performed on a simulation by simulation basis, where the user specified the scenario they wanted to investigate. Once a scenario was selected, the data were cleaned and formatted as was described before, and the clustering process began. During clustering, DBSCAN and K-means were performed on each frequency, voltage magnitude, and voltage angle data set, individually. The DBSCAN and K-means algorithms were used from a machine learning Python package, called `scikit-learn`. For that package, the original data had to be transposed, such that each row held the measurement data for each bus, and each column specified the time each measurement was recorded.

First, the transposed frequency, voltage magnitude, and voltage angle data were passed into the DBSCAN algorithm. Each data point was actually a full row of time-stamped measurements for a bus. Then the DBSCAN algorithm began by picking an arbitrary bus' measurements, then looked for `min_points` number of other bus measurements within a distance of `epsilon`. The process outlined above then continued until all data were sorted into a cluster or identified as noise.

The values of `epsilon` and `min_points` were set for each measurement type, after experimenting with what gave the best cluster results. In the future, an automatic and more sophisticated tuning approach should be used to find `epsilon` and `min_points`.

Next, K-means clustering was performed. The number of clusters formed in the DBSCAN algorithm, n , was used as an input to the K-means algorithm, along with the transposed frequency, voltage magnitude, and voltage angle data. As before, the K-means algorithm was used from the machine-learning Python package, `scikit-learn`, and followed the process described above. In our project, each column of data corresponds to measurements for each bus. However, due to the transposition required for the `scikit-learn` package, additional transposition was needed for direct comparison. Thus, the vectors \mathbf{p} and \mathbf{q} are formed from sampling rows of the transposed frequency, voltage magnitude, or voltage angle data, rather than columns as is described above. Nonetheless, the K-means algorithm functioned as expected, and sorted the original data into n clusters, with n centroids. The n centroids concisely summarize the state of the system.

1.4 Outlier Detection

In addition to clustering, the tool also performs a summary of the data to identify any buses that appear to be misbehaving more than others. To do so, the difference in measurements between time steps is calculated, and then the maximum change is saved for each bus. A histogram is created, with the x-axis designating the value of the change between time steps, and the y-axis the number of buses observing that difference.

Once the histogram is made, outlier buses are classified as the buses that fall into the last full histogram bin for frequency data, the last two histogram bins for voltage magnitude data, and the last three bins for voltage angle. As of now, the number of bins to use for identifying outliers is somewhat arbitrary, so in the future they would need to be more rigorously defined. However, for the simulations we created, the aforementioned bounds worked well.

2 Visualization Solutions

To discover insights from the PMU data, and thus enable more accurate and faster operator decision-making, visualizations were created. These visualizations sought to capture the physical footprint of the test system as well as the temporal component of the measurement data. Physical geography of the system was known from the location of substations where PMUs are actually or fictitiously installed, and temporal data from

time-stamped frequency, voltage angle, and voltage magnitude measurements.

With the advent of more data, power systems face the same challenges many industries do in making sense of all the data they have. One way to do this that has been used in power systems, as well as other industries, is to visualize it. Visualization makes it easier to understand what is happening in data by revealing patterns via colors, shapes, sizing or other graphics. In doing so, visualizations increase how well the data can be understood and how quickly. Thus, it is no surprise that industries, including the power and energy industry, are tapping into the power of visualizations to understand their data and consequently drive more effective decision-making.

The type of visualization to use depends on the application it is needed for, and this is true in power systems as well. A power system planner or an operator working on post-event analysis may prefer and require more detailed visualizations, because they have the time to carefully look through them. On the other hand, an operator who is dealing with real-time operation of the grid would want bite-sized pieces of information that they can quickly understand and act upon. They want to know as much as they can about their system (wide-area situational awareness) with as little effort as possible. Different visualizations are appropriate for each of those cases, though our research focuses on the latter one.

To date, common visualization techniques used in power systems are pie charts showing how much electrical equipment is loaded, digital numeric displays providing exact measurement values, contours showing how frequency, voltage, or phase angles vary across a geographic area, and flow arrows specifying the direction of power or the gradient of phase angles [5]. However, with the new wealth of PMU data, utilities are seeking new techniques, to take full advantage of the information PMUs can provide.

In creating our new visualizations, we sought to address as many of the challenges identified in the as possible, but we took an approach more focused on novelty and outside the box thinking in hopes of sparking ideas for future innovative visualization solutions. Our creativity was bounded only by what could realistically be implemented in a control room, mitigating as many of the issues previously addressed as possible, and following human factors' best practices.

2.1 Spatial (Geographic) Visualizations

To make the visualizations as helpful and understandable as possible, we catered to operators' spatial familiarity of their system by providing visualizations that incorporated the system's physical layout. To do this, two visualization types were created, including an animation loop of frequency, voltage, and angle contours, and a geographic plot that classified substations according to the cluster they were assigned for each measurement type.

In general, contour mapping is used for spatial visualizations, though several challenges arise for this type of visualization in power systems. First, power system data is not spatially continuous, so some values have to be assumed to correctly form a contour line or area. Second, power system data is heavily time dependent, but contours can only be created for a single static time point. Lastly, voltage contour maps have to use a per-unit system or create maps for different voltage levels, separately. Despite these challenges, one key benefit to using contour maps is that they can clearly and concisely show data for a large number of buses. Because of that important benefit, we sought to build upon the contouring visualization idea by creating animated contour maps.

The animated contour maps were created by saving snapshots of contours for frequency, voltage angle,

and voltage magnitude, in each simulation, at successive time steps. The images were then compiled into a video to create a looping animation that mimics Doppler weather radar displays. The thought behind this type of visualization is that it takes advantage of a contour's ability to concisely represent a large amount of spatially distributed information and allows for the time component to be captured via time stepped changing contours. In an operational setting, the animation loop would continue to be refreshed with new data, though data before a certain interval, say 60 minutes, would be discarded. An operator can pause or rewind the loop, if they see something developing in their system.

Representing time by motion (animation) requires constant monitoring of the display for an operator to perceive the development of a grid event in time. Looping helps preserve the event in time, relieving operators from such vigilance requirement. Nonetheless, watching an animation loop still requires more time than reading a static graph or tabulated data [6]. Sample animation loops can be viewed at [7].

The other approach we used to represent the geographical component of PMU data was to show the location of buses (spatial data) with markers on a map, and then color them according to which cluster they fell into, derived from measurement data (temporal data). Data points were shaped according to the nominal voltage level of each bus. The idea behind this was to visually show regions of the grid that behaved in a similar way. Figure 1 shows a plot of buses being monitored by PMUs, whose cluster and color assignments were determined from frequency simulation data that was sent through the clustering process described above. By using individual data points, rather than a contour, the requirement for spatial continuity was circumvented. Additionally, because individual buses were plotted, an operator may be able to more quickly identify the exact set of problematic buses, rather than just a general area of concern.

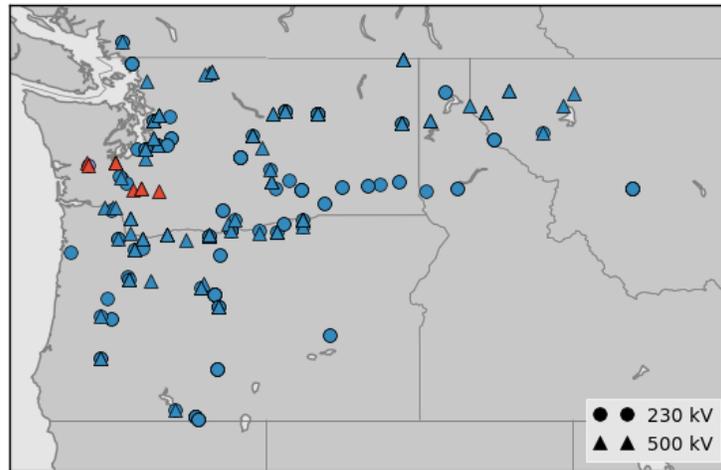


Figure 1. Geographic plot of bus locations. Marker colors specify the cluster to which each bus belongs, and the shape the voltage level each bus operates at.

In the visualization of Figure 1, and all visualizations we created, extraneous lines, colors, and data were excluded to focus the attention of the operator on the most important information, for example, by using gray background colors rather than topographical colors. This also ensured that more colors were available to represent and emphasize system data.

2.2 Temporal Visualizations

Temporal data is often expressed with line plots, because they easily display how data changes for each time step plotted on the x-axis. However, using line plots for tens or hundreds of buses quickly makes the plotted data cluttered and unhelpful. Because of this challenge, two techniques were used to limit the amount of data presented in line plots. First, clustering was used to reduce the data down to representative cluster centroids, which are mean values of the data in a cluster. Second, a drill-down technique was used, such that data was only shown for what is selected by the user.

In Figure 2 line plots are shown for each bus being monitored by PMUs (left) and the cluster centroids (right), for each measurement type. It is clear that when all PMU data are plotted the result is a cluttered display of overlapping lines that do not provide very useful information. On the other hand, when just the cluster centroids are plotted (right side of Figure 2), they are able to capture the envelope response of the system for each measurement type, and in so doing provide a concise summary of the state of the system using only a few lines, rather than hundreds. Unfortunately, using the clustered data does not provide detailed information about a specific bus or substation, but it does offer a summary that would allow an operator to narrow down the area they should look at next in their investigation.

Figure 3 provides a visual of a more granular view of the clustered measurement data, by plotting the original measurement data, within each cluster, alongside the cluster's centroid data points. The data shown in Figure 3 is for voltage magnitude clusters, though the same sort of plots could be formed for frequency or voltage angle. As before, you can see that the cluster centroid, for the most part, captures the general response of the system. The cluster centroid shows where events occurred via the sudden spikes or divots, and in so doing provides focus for further investigation by an operator.

Misbehaving outlier buses can be identified for the data as a whole or for each individual cluster. This extra information can supplement the cluster centroid data to provide both a big picture understanding of the system, with the cluster centroids, and a more granular view of the system, with identification of outlier buses. The outlier buses are the ones that most likely need to be dealt with to mitigate power system problems. Figure 4 provides a visualization of outlier buses for cluster 4, shown in Figure 3. A display like Figure 3 could be used in combination with other temporal displays.

In an operational setting, the plots in Figures 1 to 4, would not, by themselves, be helpful. Thus, the next section describes our ideas for how elements from the aforementioned spatial and temporal plots can be combined to provide improved wide area situational awareness and enough detailed information to help an operator pin point a problem and fix it.

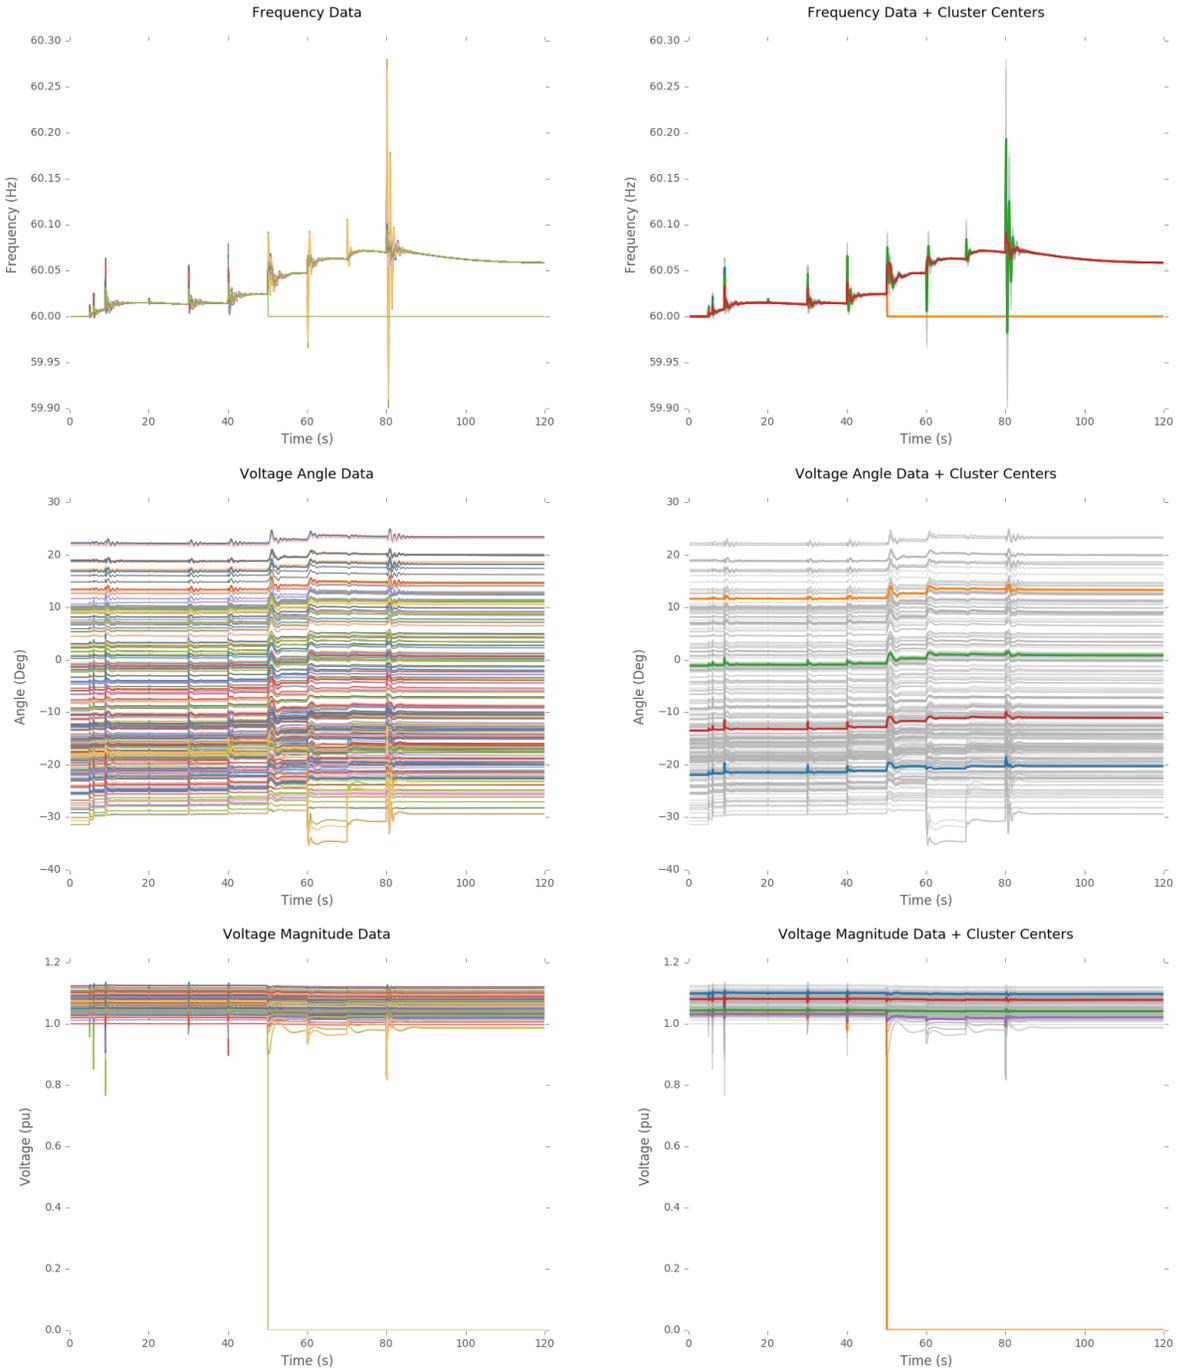


Figure 2. Temporal line plots of PMU data, for all buses being monitored (left) and summative cluster centroid line plots (right) for each measurement type.

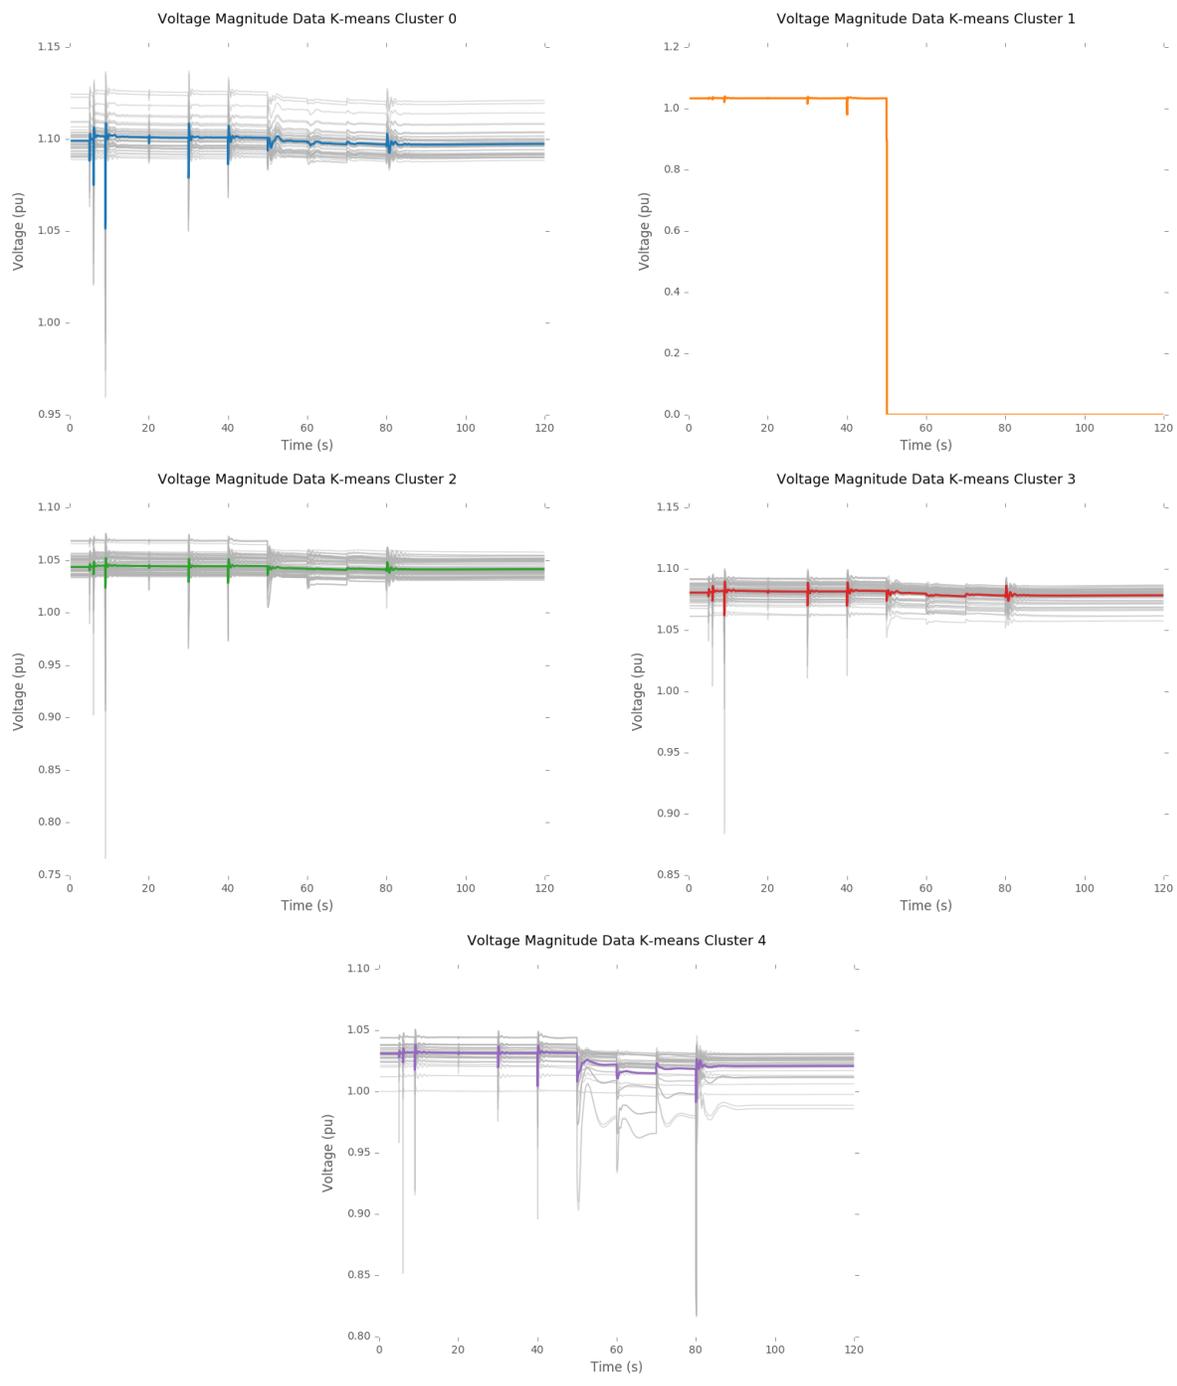


Figure 3. Temporal line plots of PMU data, for all buses being monitored (left) and summative cluster centroid line plots (right) for each measurement type.

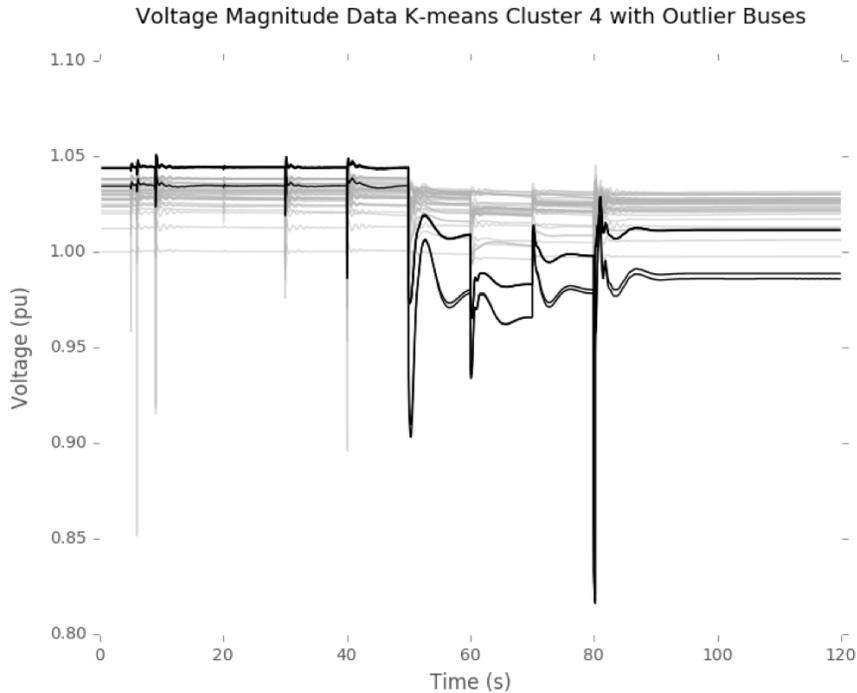
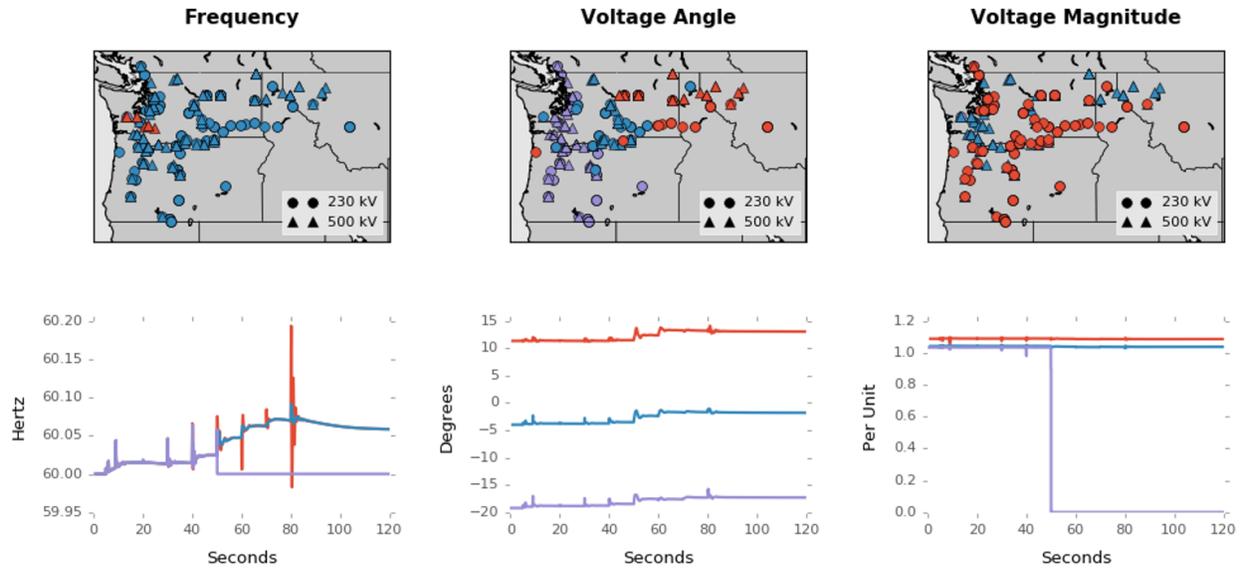


Figure 4. Line plots of data for a cluster of voltage magnitude PMU data, with outlier buses identified by the black dark lines. Substation names for the outlier buses can easily be provided in the plot, but were kept hidden for confidentiality.

2.3 Visualizations to Pinpoint Problems

Two interactive visualization panes were created to give operators better situation awareness (SA) and the ability to quickly identify problems. Each display uses visualizations that capture both the temporal and spatial elements of PMU data. The first interactive display shown in Figure 5, is made from data after it has been clustered using the process described above. The plots in the top row show the geographic location of each bus in BPA's system. The markers are shaped according to their voltage level and colored according to the cluster they have been assigned. The second row of plots shows the cluster centroids, to summarize the overall response of the system, without plotting data from every bus. This ensures a concise, but complete, picture of a large electric grid. The last blank row of plots is reserved to show measurement data for buses that are selected from one of the geographic plots. For example, the user can draw a rectangle over a section of one of the plots in the first row, and then the data for the buses contained in the rectangle will be plotted in the last row. The user selection process is shown in Figure 6. Note that the colors in each column of plots are associated, though not necessarily the colors in plots of the same row.

There are several benefits for using this type of display. First, there are only as many colors used as there are clusters, which ensures the operator can focus on what the data are saying rather than how they are saying it through colors. Second, a concise and summative state of the system is provided by using cluster centroids, rather than plotting data for every bus in the network, helps reduce clutter. However, if a more detailed view is desired, the display also allows for that, through user selection of buses on the geographic plots. This balances the need for a big picture perspective and more granular details, with only one mouse click. Finally, because this display incorporates the geographic locations of buses, outside information like



No plots appear in this row until the user makes a selection of buses they want to zoom in on, in one of the geographic maps.

Figure 5. Display of (top) geographic locations of buses and the clusters they belong to and (bottom) the cluster centroids, for the earthquake scenario.

weather radar could be overlaid on the plots to more quickly diagnose the cause of a problem, especially for weather related incidents.

The second interactive display is shown in Figure 7. The basis for the visualizations in this display is the underlying grouping of data based on the maximum differences between time steps for each bus, described above. Histograms are provided in the first row and show the maximum difference between adjacent time steps, for each bus. Sections of each bin are colored according to what time window the maximum difference is observed, for each bus. The second row of plots shows the measurement data for outlier buses contained in the last one or two full bins of the histogram above them. One bin is used for frequency and two bins for voltage angle and magnitude. The color of each line does not correspond to the histogram, but is uniquely colored to distinguish between various buses. The final row of plots shows the geographic locations of each outlier bus identified for each measurement type—frequency, voltage angle, and voltage magnitude. Note that the colors between plots are not necessarily associated.

Just like the previous display, this display is also interactive. A user can select a bin in one of the histograms that they want to investigate further. By doing so, the middle row of line plots updates to reflect the data for the buses contained in the bin the user just selected, as does the geographic plot to show the correct buses that were selected. The new line plot also includes the corresponding bus number for the data, which allows an operator to quickly identify the problematic buses. Figure 8 shows the updated display, from Figure 7, after a user selected to view the very last bin of the frequency histogram. The two buses contained in that bin then have their data displayed in the second row of line plots and their location in the third row of geographic plots.

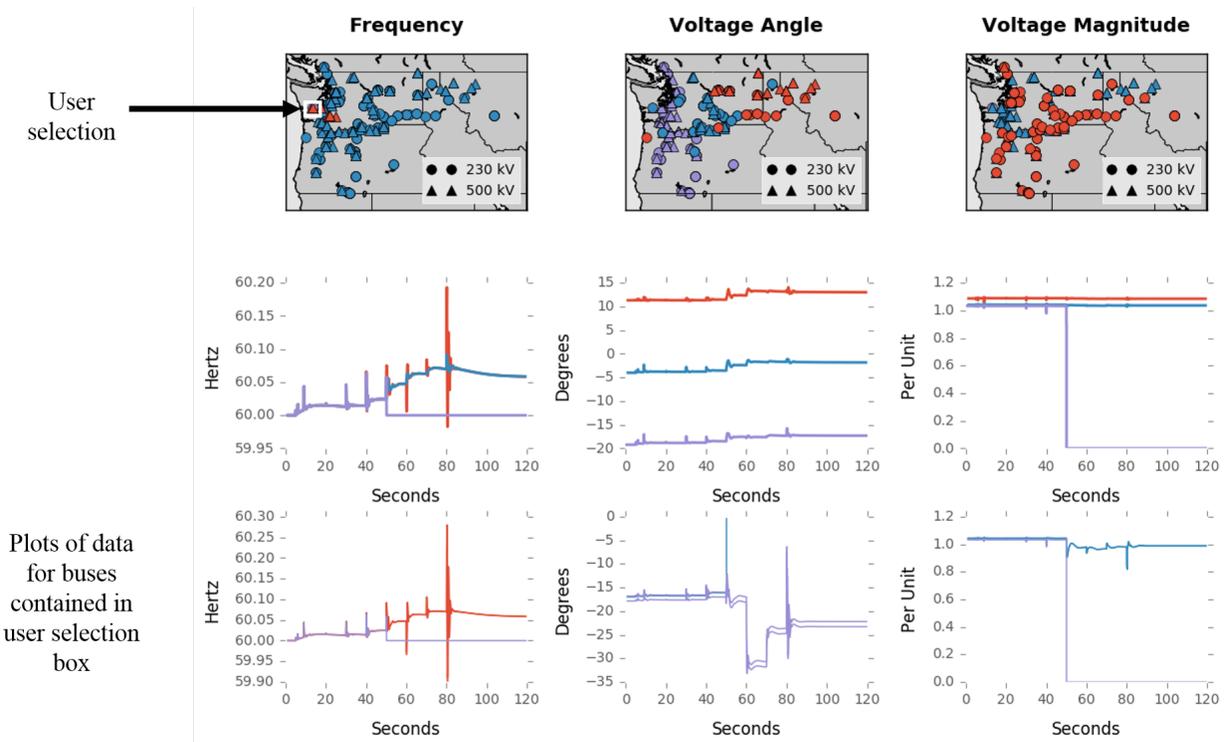


Figure 6. Display of (top) geographic locations of buses and the clusters they belong to, (middle) the cluster centroids, and (bottom) the data for the buses selected by a user-drawn rectangle in the frequency geographic plot, for the earthquake scenario. Colors are associated in plots within each display column, not necessarily the row.

This type of display makes it easy to concisely summarize drastic changes in measurement values and to show in which 10-s window those changes are observed. The display allows an operator to quickly identify outlier buses and see the temporal nature of the data for those buses, as well as their geographic footprint. By doing so, third party information, like weather radar, can be incorporated into the display to help diagnose problems and then correct them. One problem is that there may be too many colors used in this display, which could distract or strain the operator from focusing on what is most important.

The difference in displays can be characterized by function, where the first display allows for interpretation and search based on geography, and the second allows for interpretation and search based on the most drastic changes between data points. As mentioned before, looping contour videos may provide a nice addition to each display.

The displays shown in Figures 5 to 8 are for the earthquake scenario, occurring on the western coast of Washington. In the first display, the blue colored markers of the frequency geography plot seem to indicate buses that are either the cause of a problem or are affected by one. In fact, the buses shown in blue are the buses that were manipulated in the simulation and/or those that are electrically close to the manipulated buses. In the second display, the largest deviations are observed in the region where the earthquake's damage was greatest. Similarly, the displays of Figures 9 and 10 correctly identify buses that were directly affected

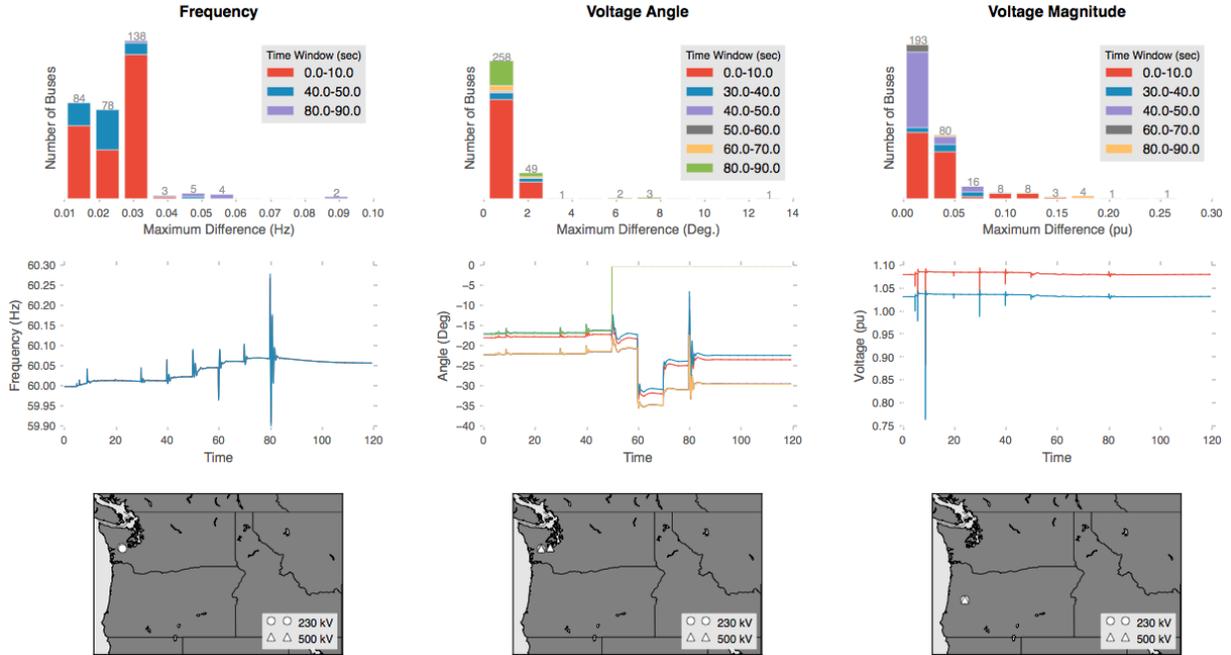


Figure 7. Display of (top) maximum changes between time steps observed for each measurement type, (middle) measurement data for each bus contained in the last bin, for frequency, and last two bins for voltage angle and voltage magnitude, and (bottom) geographic location of the buses identified as outlier buses in the last one or two bins, for the earthquake scenario.

by an ice storm or the buses that are electrically close to those buses.

All in all, the displays discussed in this part of our report are meant to inspire new ideas for visualizations, and are not necessarily meant to be final products. They may be beneficial in their own right for real-time operation of the grid, or could help in offline studies of events to help operators come up with new procedures, called remedial action schemes, that dictate what actions an operator should take when a certain event is observed. Indirectly, these visualizations would then help in future real-time operation of the electric grid.

3 Summary and Discussion

In the area of visualization, we sought to think “outside the box” and spent many hours poring over visualizations of every shape and variety. We saw many beautiful visualizations that sparked ideas for our own research, but ultimately concluded, as others have, that visualizing spatio-temporal data is quite challenging, especially when the visualizations have to be interpreted quickly.

We noticed that previous research in power system visualization has largely been done in silos, where experts in power systems, statistics, human factors, or design have yet to come together to collaboratively tackle visualization problems. For this project, we changed that paradigm and brought together experts from both power systems and human factors. In doing so, we tapped into the power of collaboration and a diversity of skill sets. Though advantageous for creating impactful visualizations, it added a layer of complexity

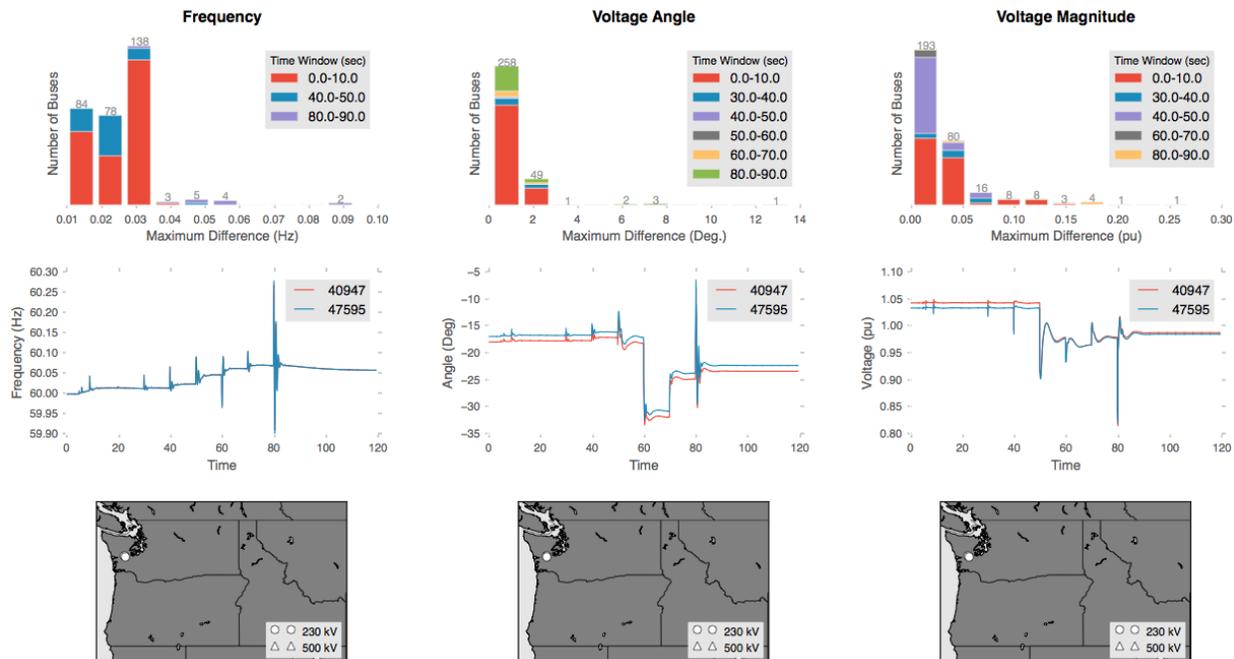


Figure 8. Display of (top) maximum changes between time steps observed for each measurement type, (middle) measurement data for the buses contained in the last bin, as selected by the user from the frequency diagram, and (bottom) geographic location of the two buses contained in the last bin of the frequency histogram, for the earthquake scenario.

in trying to create data suitable for the human factors team to deal with and which was understandable. To ease that burden, we created a software tool to standardize and automate the process of transforming power system data into a well defined format for which visualization prototypes could quickly be developed. The idea was to create a bridge between power system data and data visualizations, so that visualization experts can focus on creating visualizations rather than waste time trying to create data and understand power system intricacies. As long as data input to our program is of the form given in the appendix, and visualizations are created in Python, the tool should be able to get the user from file selection to visualization within a few clicks and a few seconds.

After exploring various clustering algorithms, we ultimately decided to use K-means and DBSCAN because of their speed and accuracy. Beyond that, DBSCAN was chosen for two reasons. First, because one of the outputs of the DBSCAN algorithm is the number of clusters that can be formed, we used that to initialize the number of clusters used in the K-means clustering step. Second, we thought the input parameters to DBSCAN, epsilon and min_samples, could be helpful for utilities to define based on what made sense for their system. As such, the clustering program could better cater to a utility, and thus make better summative representations of power system data. However, if the user does not want to specify either the number of clusters to use for K-means or the input parameters for DBSCAN, an alternate method can be used. This method runs the K-means algorithm for n number of clusters, where n is cycled from 2 to 10. A score representing how well the data fits in each cluster is created for each run, and then the program selects the number of clusters, n, which has the lowest score, or at least the first minimum score. Although this approach eliminates the need for any input parameters, it is slower since the K-means algorithm has to be run at least two times in order to generate clustered data. In the software tool, a user can select whether they

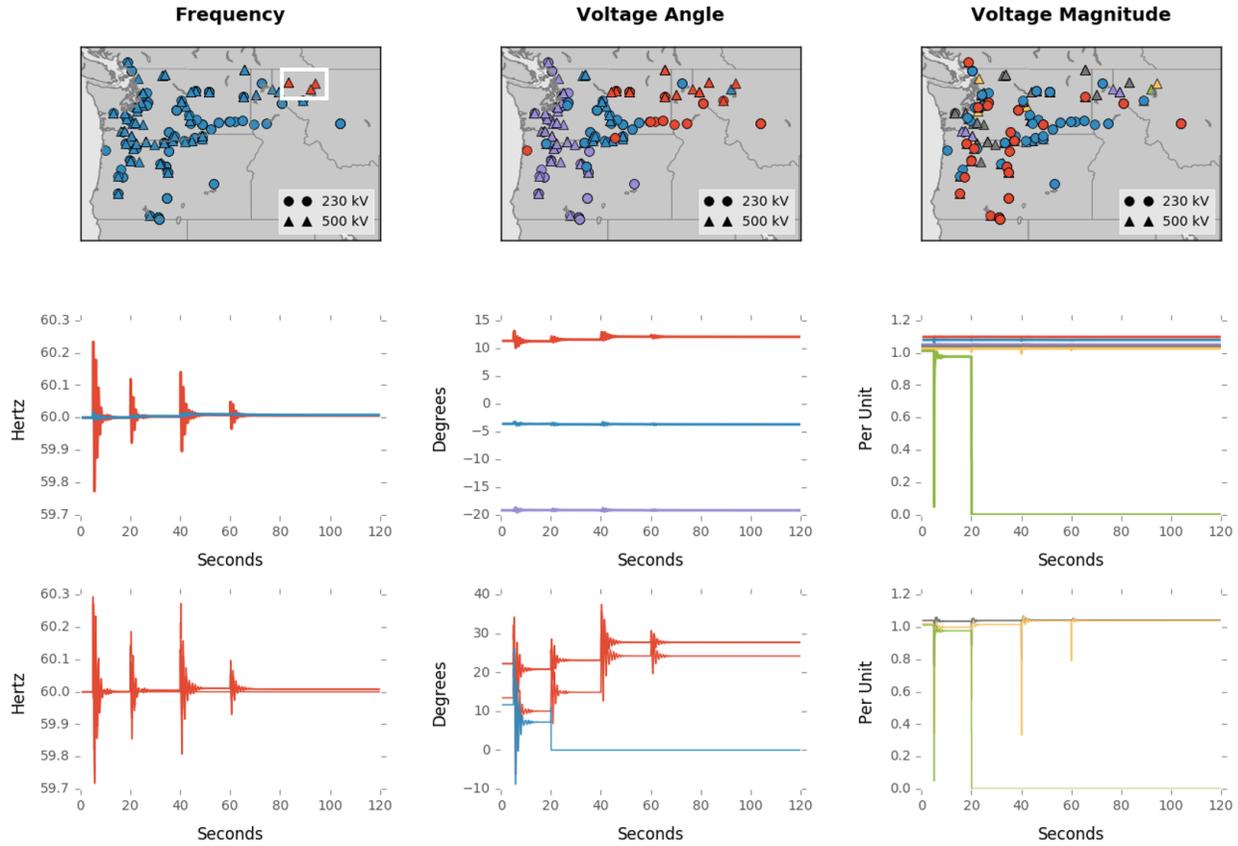


Figure 9. Display of (top) geographic locations of buses and the clusters they belong to, (middle) the cluster centroids, and (bottom) the data for the buses selected by a user drawn rectangle in the frequency geographic plot, for the ice storm scenario. Colors are associated in plots within each display column, not necessarily the row.

want to use DBSCAN or the K-means scoring method for initializing the K-means algorithm.

Finally, by creating visualizations with geographic maps, non-power system data can easily be integrated to help operators diagnose problems in their systems. For example, adding weather radar to one of the geographic maps in Figures 5 to 10 could enable an operator to see if a storm is causing, or could cause, an event in their power system, so that they can take the appropriate actions. Without PMU data and its geographic information, these sorts of displays would not be possible and the actions an operator should take less obvious.

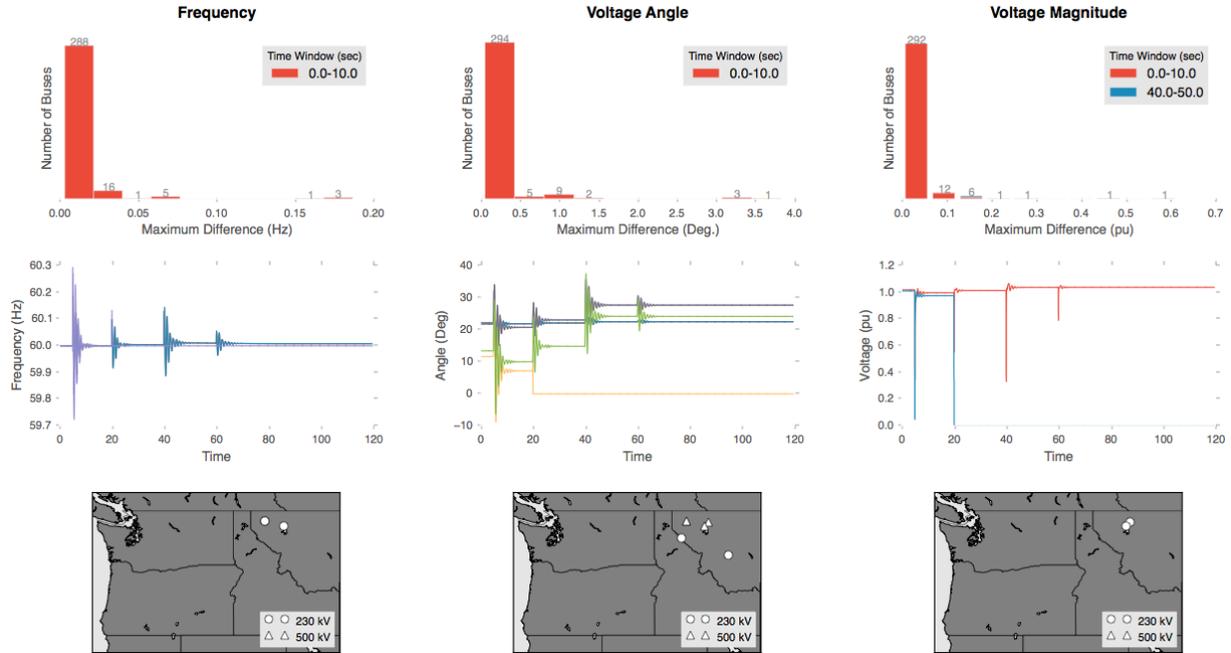


Figure 10. Display of (top) maximum changes between time steps observed for each measurement type, (middle) measurement data for each bus contained in the last bin, for frequency, and last two bins for voltage angle and voltage magnitude, and (bottom) geographic location of the buses identified as outlier buses in the last one or two bins, for the ice storm scenario.

References

- [1] A. Von Meier. Occupational cultures as a challenge to technological innovation. *IEEE Transactions on Engineering Management*, 46(1):101–114, 1999.
- [2] N. Harris. Visualizing DBSCAN Clustering. <https://www.naftaliharris.com/blog/visualizing-dbscan-clustering/>, 2015.
- [3] M. Ester, H-P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD96)*, volume 96, pages 226–231, 1996.
- [4] Scikit-learn.org. Scikit-Learn 0.18.1 Documentation – 2.3.2 K-means. <http://scikit-learn.org/stable/modules/clustering.html#k-means>, 2016.
- [5] T. J. Overbye. Visualization enhancements for power system situational assessment. In *Proceedings of the IEEE Power and Energy Society General Meeting–Conversion and Delivery of Electrical Energy in the 21st Century*, pages 1–4. IEEE, 2008.
- [6] K. M. Gegner, T. J. Overbye, K. S. Shetye, and J. D. Weber. Visualization of power system wide-area, time varying information. In *Proceedings of IEEE Power and Energy Conference at Illinois (PECI)*, pages 1–4. IEEE, 2016.
- [7] PowerWorld Corporation. Animated transient stability contingencies in simulator version 19. <https://www.powerworld.com/animated-transient-stability-contingencies-in-v19>.

Part III

Visualization of Multiple Time Series

1 Statistical Background

We considered different approaches to the visualization of a collection of time series from the simulated grid events. Our ultimate goal was to visually detect any existing anomaly/novelty in the time series worthy of being flagged for examination or even immediate action by an operator monitoring the grid of electric power stations. To gain insights into some of the potential models underlying the data, we take a microscopic look at a handful of scenarios, and specifically find out which (if any) of the measurements reveals more in the way of detecting novelties and anomalies.

The stacked plot shown in Figure 1 is an exemplar of the many time series data from the simulated scenarios explored and scrutinized in this project. As can be seen, voltage and phase angle appear to be different in dynamics from frequency. While there appears to be occasional large deviations in the frequency measurements, voltage and phase angle both seem to be uneventful judging by the unchanging dynamics. In this sense, Frequency appears to have a greater potential for helping detect novelties and anomalies in the electric grid. A thorough study of the dynamics of each of the three measurements allows the detection of some of the changes that were difficult to spot in the raw series.

More formally, we use $X_{s,t}$ to denote observation of bus s at time t . The observation could be the frequency, the phase angle or the voltage or any other of the available variables from the grid. We assume here that there are S buses in total, all producing observations. Now, if we observe from time $t = 1$ to $t = T$, the raw data available can be represented as a matrix \mathbf{X} given by

$$\mathbf{X} = \begin{bmatrix} X_{1,1} & X_{1,2} & \cdots & X_{1,t} & \cdots & X_{1,T-1} & X_{1,T} \\ X_{2,1} & X_{2,2} & \cdots & X_{2,t} & \cdots & X_{2,T-1} & X_{2,T} \\ X_{3,1} & X_{3,2} & \cdots & X_{3,t} & \cdots & X_{3,T-1} & X_{3,T} \\ \vdots & \ddots & \cdots & \ddots & \cdots & \ddots & \vdots \\ X_{s,1} & X_{s,2} & \cdots & X_{s,t} & \cdots & X_{s,T-1} & X_{s,T} \\ \vdots & \ddots & \cdots & \ddots & \cdots & \ddots & \vdots \\ X_{S-1,1} & X_{S-1,2} & \cdots & X_{S-1,t} & \cdots & X_{S-1,T-1} & X_{S-1,T} \\ X_{S,1} & X_{S,2} & \cdots & X_{S,t} & \cdots & X_{S,T-1} & X_{S,T} \end{bmatrix}$$

Of great interest in the analysis is the vector $X_s = (X_{s,1}, \cdots, X_{s,t}, \cdots, X_{s,T})$ representing the realized time series of observations for bus s . In other analyses, the focus shifts to $X_t = (X_{1,t}, \cdots, X_{s,t}, \cdots, X_{S,t})$, which is the S -dimensional vector at time t for all the S buses.

When we consider the state of the grid at a given time point t and focus on the measurements coming from each bus, it is reasonable and desirable that the repeated measurements within each bus will be very similar in values or almost nearly equal when the grid is “healthy”. For this reason, fluctuations are expected to be rare with the consequence being that the variance in the vector $(X_{1,t}, X_{2,t}, \cdots, X_{S,t})$ is typically tiny.

Our exploration reveals that signal detection can be hampered by the minuteness (smallness) of the variations (fluctuations) unless we pre-process the data. For that reason, we explored various preprocessing

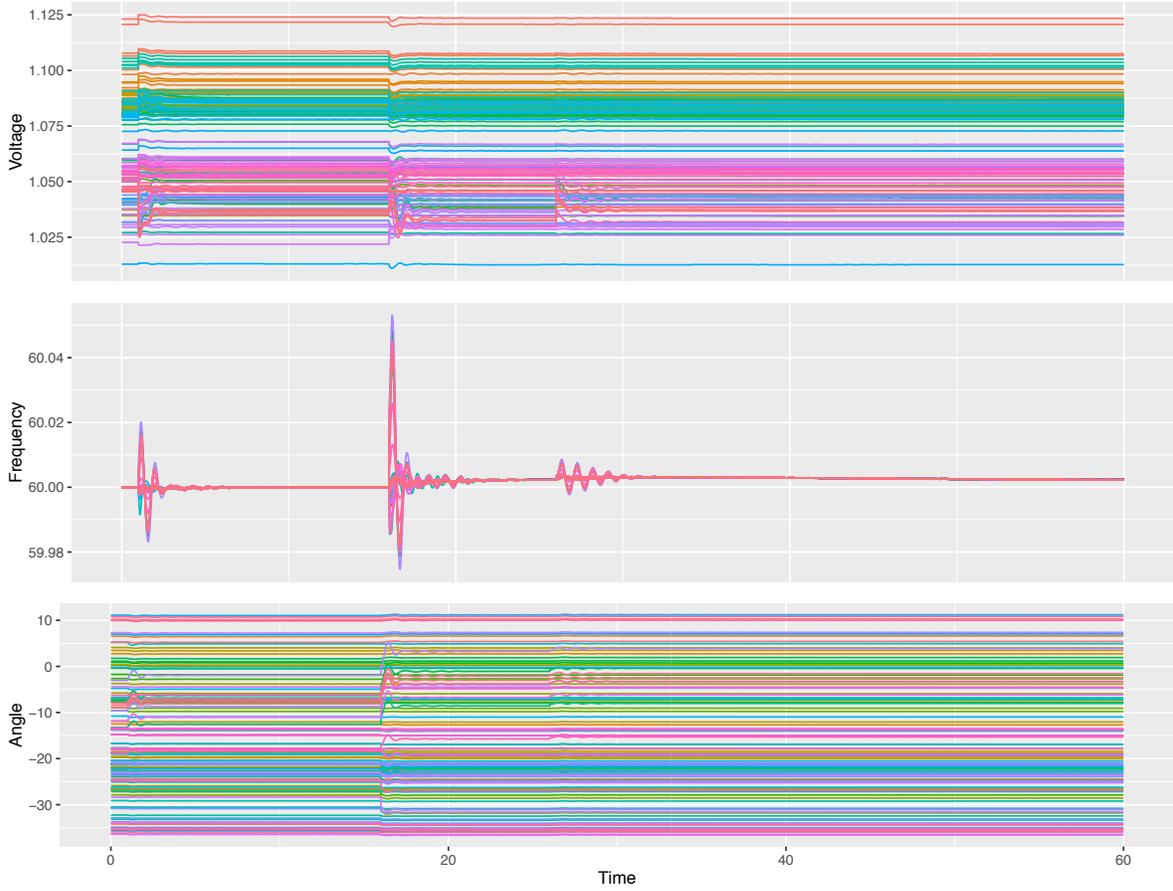


Figure 1. Raw series on voltage, frequency and phase angle measurements for the McNary attack scenario. Note how the variations in voltage and phase angles between the buses make for very cluttered plots.

schemes to help isolate and study only the variations of interest. The first is standardization

$$X_{s,t} \mapsto \frac{X_{s,t} - \bar{X}_t}{\sqrt{\frac{1}{S-1} \sum_{s'=1}^S (X_{s',t} - \bar{X}_t)^2}} \quad (1.1)$$

It turns out also that several time points yield a variance of zero, leading to a numerical problem that we circumvent by using the internal zero of the computer, so that our mapping is

$$X_{s,t} \mapsto \frac{X_{s,t} - \bar{X}_t}{\sqrt{\frac{1}{S-1} \sum_{s'=1}^S (X_{s',t} - \bar{X}_t)^2 + \epsilon}} \quad (1.2)$$

where $\epsilon > 0$ is the so-called "zero" of the compiler being used. Another pre-processing explored is the unit cube or cubitization, namely

$$X_{s,t} \mapsto \frac{X_{s,t} - \min_{s \in [S]} X_{s,t}}{\max_{s \in [S]} X_{s,t} - \min_{s \in [S]} X_{s,t} + \epsilon} \quad (1.3)$$

As we saw with Figure 1, it is hard to tell anything consequential about voltage and phase angle variability because of the sheer flatness of the corresponding series. It is crucial to note that their flatness does not necessarily mean that all is well (constancy). It could be hiding minor fluctuations that then grow and later become a major problem. The raw series in most cases may not reveal such situations. Therefore it is imperative to perform some statistical modeling on the series to hopefully extract some intrinsic patterns of behavior. Ideally, one would like to monitor and model the evolution of the time series well enough to

1. Detect anomaly (novelty) worth of being acted upon by the operator, or even better
2. Extract deep knowledge from the time series, so as to accurately predict future potential anomalies, and therefore warn the operator to be prepared for preemptive action to prevent catastrophes
3. Perform an *a posteriori* analysis of the time series after a catastrophe has happened, with the posterior analysis aimed at comprehending the causes so as to prevent/avoid them in the future

Prior to diving deep into full scale modeling however, it makes sense to explore simple and easily computable derivatives of the series like first order differences, first order rate of change and even second order difference. The first order rate of change for instance, because it is a proportion, can be very revealing when it comes to wild changes. For instance, a change of more than 50%, however tiny the magnitude might be, should be taken seriously. To further explore this idea, we generated stacked plots for the three measurements, each stacked plot featuring the raw series along with the derivatives mentioned earlier. For clarity, we have the following:

1. The raw time series itself, namely $X_{s,t}$, for $t = 1, \dots, T$. Various computations can be performed on this raw data to extract features worthy of further examination. For instance, one could perform an anomaly detection on this series using basic outlier detection methods for univariate data.
2. The first order difference $Y_{s,t} = X_{s,t} - X_{s,t-1}$, for $t = 2, \dots, T$. Clearly, $Y_{s,1} = 0$, since there is no change at the initial time point. This provides the most basic tool to probe the dynamics of the series with the potential of detecting any rapid changes worthy of further characterization
3. The first order rate of change $Z_{s,t} = Y_{s,t}/X_{s,t-1}$, $t = 2, 3, \dots, T$. We must also have $Z_{s,1} = 0$ since there is no change at the initial time point. By removing the scale from the first order difference, the rate of change provides a very good unit-free measure of the dynamics of the series.
4. With the first order difference revealing aspects of the velocity, it is interesting to also consider the second order difference which then tells us about the acceleration. In an effort to discover the inner workings of the stations, the acceleration could help in early warning and forecasting.
5. Just as the rate of change was considered in the first order, a second order rate of change would help grasp the percentage by which the acceleration has changed from one time point to the next.

Plots such as the example in Figure 2 were generated to further explore this idea around the potential of derivatives to reveal novelties and anomalies of interest to the operator. In each case, both the raw series and its derivatives are plotted together in a stacked format, to help detect whatever might be interesting in the data, in the sense of anything requiring a reaction/response by an operator monitoring the grid of stations. It is crucial in all these cases to be mindful of *false alarms*, and ideally detect and reveal only the truly action-worthy novelties/anomalies. All the above mean that any visualization produced must be the visible representation/rendering of a robust novelty detection mechanism applied to the data.

However, visualizations such as the line-plot in Figure 2 can be potentially very useful to operators monitoring the grid. As can be seen in the figure, between-buses variability makes the display of raw time series

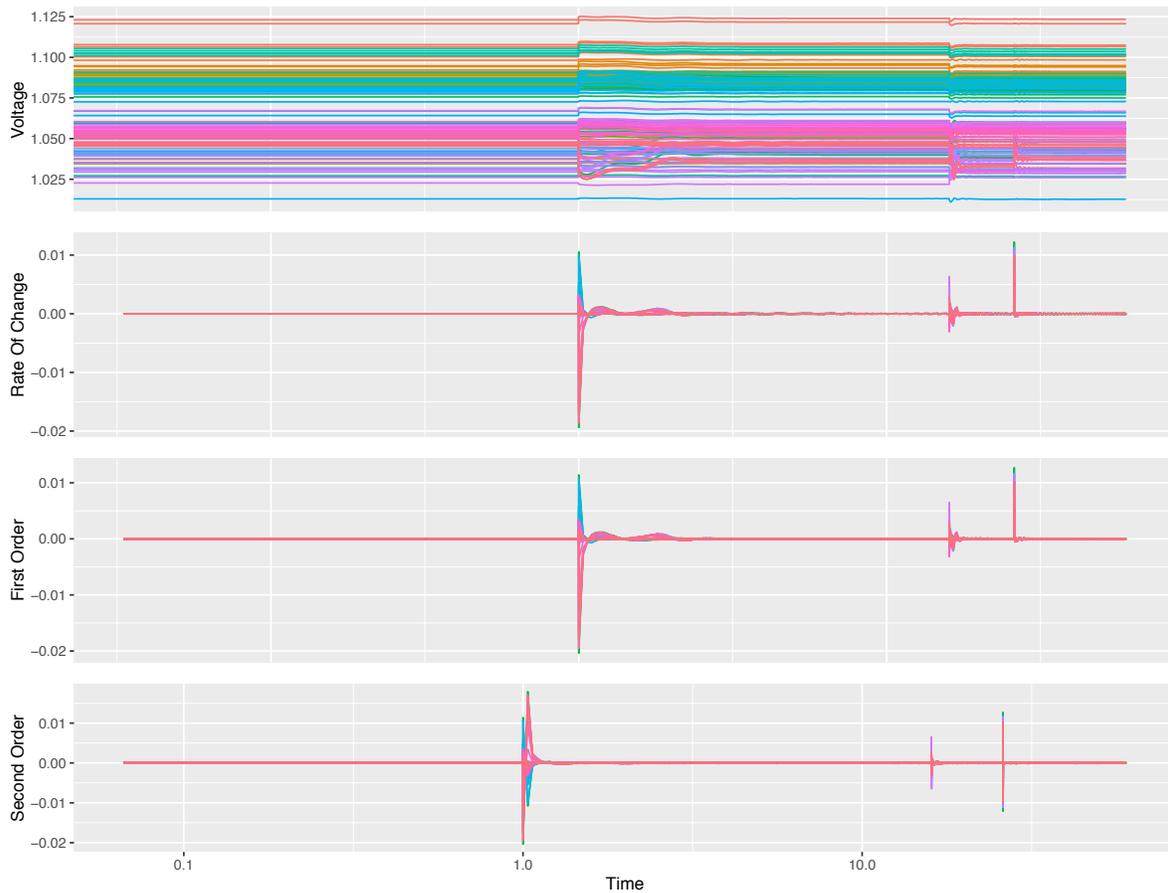


Figure 2. The McNary attack scenario, raw series and derivatives of voltage measurements.

largely unusable. A fortuitous solution was offered by calculating the first-order rate of change of the buses. This transformation, in addition to providing operators with additional information (i.e., the rate of change in selected variables) and facilitating Level 3 SA (projection of system state into the future) also would isolate the affected buses and greatly reduce the amount of information on the display.

2 Prototype Display Design

In the simulations performed the number of buses with operationally significant changes in the measured variables is small enough to be shown on a single display and identified by labels or a legend. Other key design criteria were to display the three most relevant variables (phase angles, frequency, and voltage) in a single display, against a common timeline. Furthermore, it was determined that at least 5 minutes' worth of data should be viewable at any time, the size of the display should be compact enough for placement in the video-wall, and enough detail of the event should be visible to allow for accurate diagnosis. These criteria were met by a logarithmic transformation of the time (x) -axis so that most recent data are displayed on a "stretched" scale and historical data are increasingly compressed the farther into the past the scale reaches.

In this prototype display the three key design features are (1) visual isolation of “misbehaving” buses by displaying first-order differences or rates of change, (2) presentation of the three most critical variables on a common time axis so that their interactions may be easily perceived, and (3) “compression” of the reverse logarithmic time axis so that as much history (to the left of the axis) as possible may be viewed at all times while most recent events to the right of the axis are displayed with high resolution. These features are shown in Figure 3.

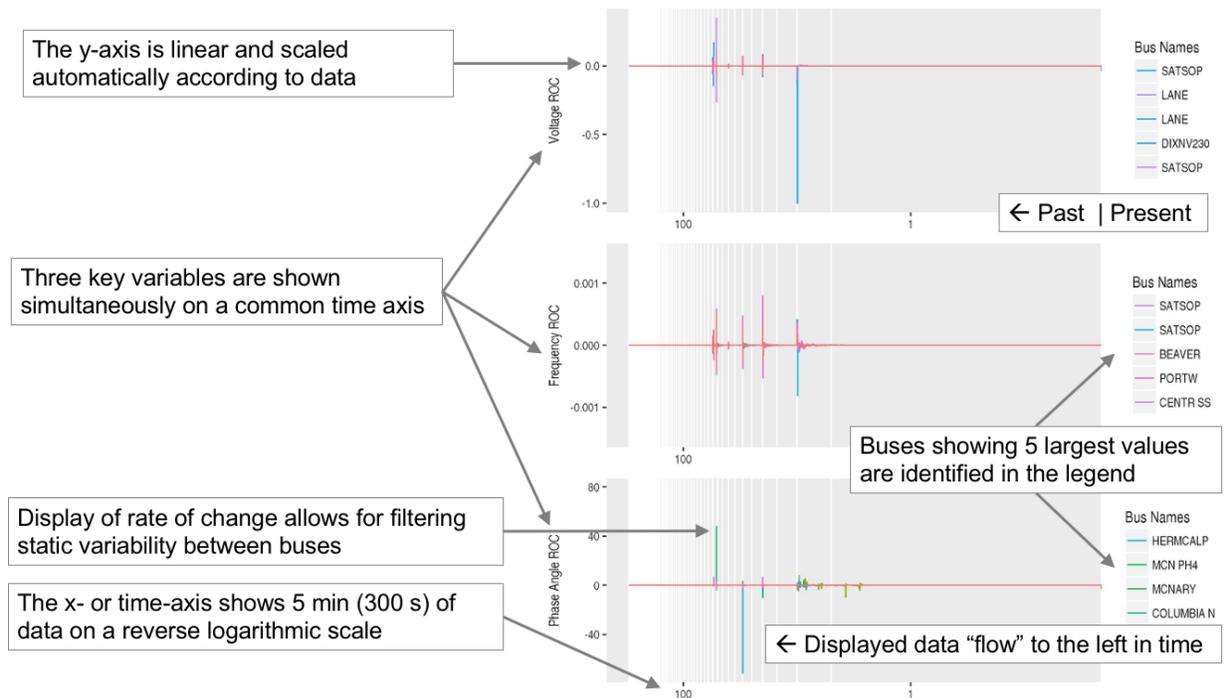


Figure 3. The earthquake scenario shown on a prototype display, featuring rates of change of the three key variables (voltage, frequency, and phase angles) plotted against a common reverse logarithmic time (x) -axis.

In Figure 4 are all the simulated scenarios as they would appear in the prototype display. Although the primary purpose of the display is to alert the operators to a changing state of the buses (rate of change is shown), it is possible that different events would also present different patterns that may be recognizable by experienced operators. In other words, our hope is that the (mostly) primary data (one derivative, first-order difference, or rate of change) would allow experienced operators have insights into the events on the bus level, and thus some diagnostic power that would guide their actions corresponding to the event displayed. Of course, such capability should be tested with dynamic displays and many more scenarios that the 8 we simulated.

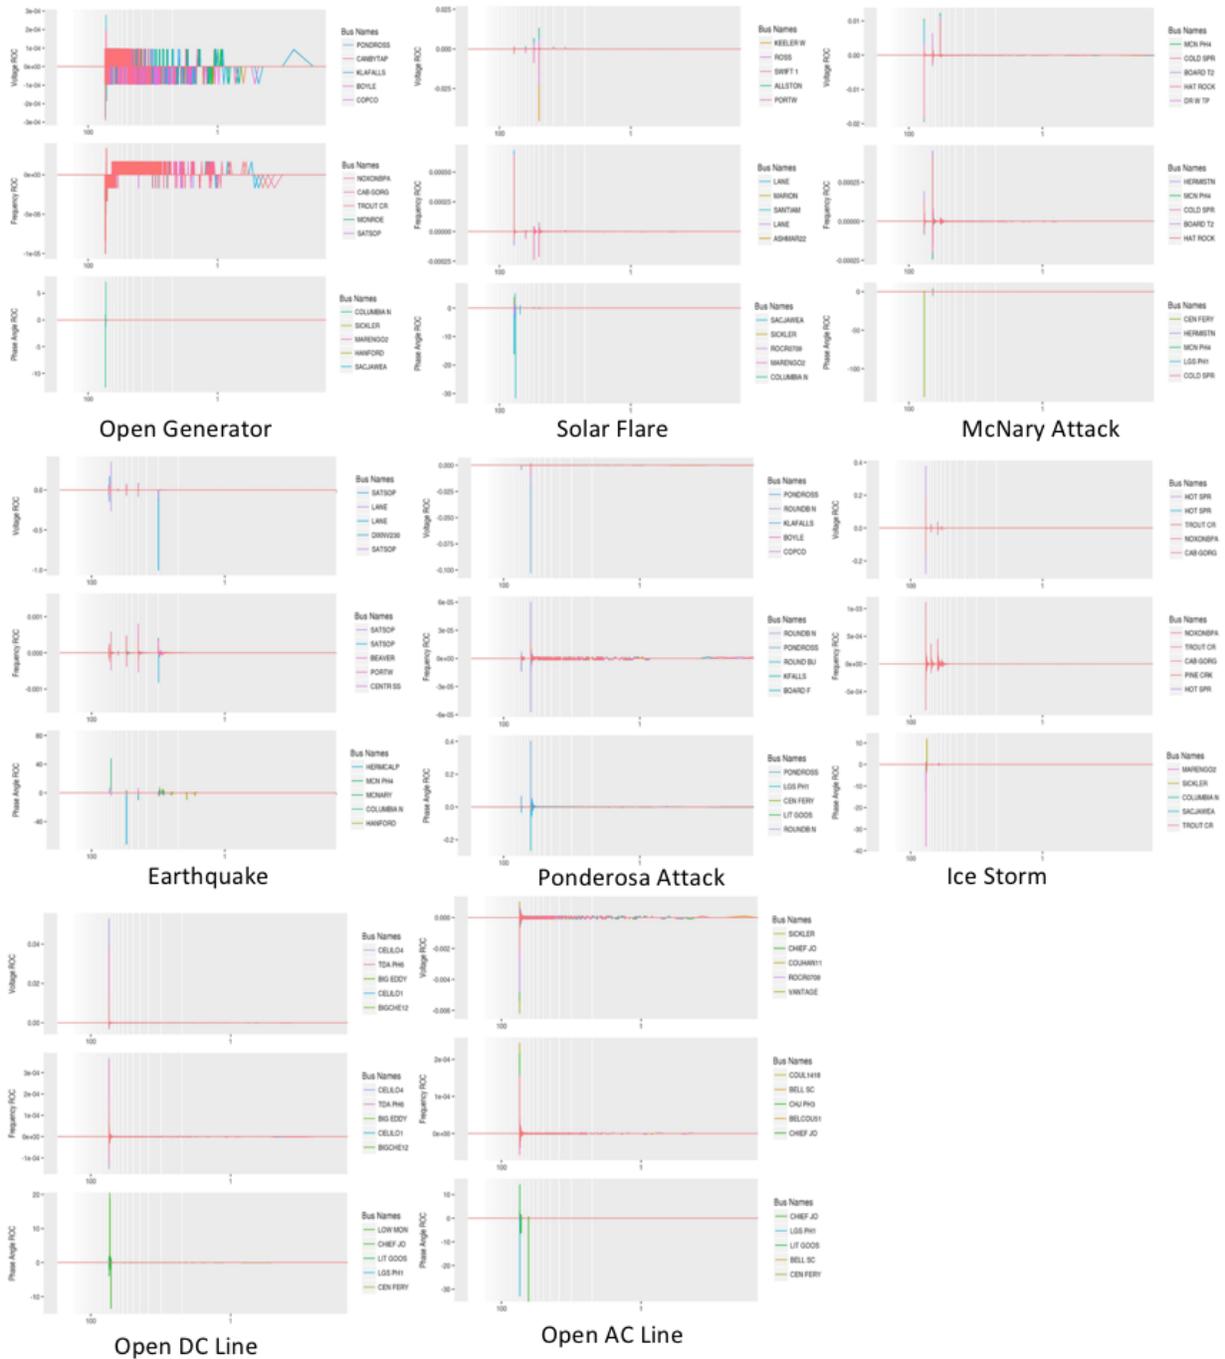


Figure 4. The different simulated scenarios depicted on the prototype line plot displays.

2.1 Implementation in R

The visualizations depicted in Figures 3 and 4 were created in R, a free software environment for statistical computing and especially for graphics, which fit our needs very well. The following describes the steps and R libraries and packages used to create the visualization:

1. **Calculating Rate of Change:** We use the library TTR provided by R. This package allows us to construct technical trading rules (TTR) with R. It is used to calculate the Rate of change over n periods using the function ROC. We give the input data and the type of the data. In our case the type of the data is discrete.
2. **Scaling the axis:** We use different properties of of the scales when we plot the data. To enable this we use R's scales library. The scales package is made up of the following interdependent components for aesthetics, transformations, bounds, breaks and ranges. We use these properties while plotting the axis is reverse logarithmic scale.
3. **Data selection and manipulation:** While dealing with data frames we use R's tidyr and dplyr libraries. Both of these packages helps in easier manipulation and the selection of the rows and columns along with their data frames.
4. **Data transformations for the plots:** Before we give the data to plot, plotting library requires the data in an appropriate form. In our case we use R's reshape2 library. This package's melt() function is used so that it takes the data in wide format and stacks a set of columns into a single column of data. To make use of the function we need to specify a data frame which is frequency, voltage and the phase angle. The id variables (which will be left at their settings) is the time and the measured variables (columns of data) to be stacked. The default assumption on measured variables is that it is all columns that are not specified as id variables.
5. **Plot:** To plot the data we use R's ggplot2 package. It is a system for "declaratively" creating graphics, based on [1].
6. **Stacking the plots:** The individual plots are stacked using R's ggpubr library. This package allows to align multiple ggplots into single stack. We did this by employing the ggrrange() function.

3 Detection and Display of Anomalies in the Data

The most striking insight gained from plotting the raw series along with its derivative is that even when the raw series is unclear as to at which time point the operator might need to pay attention, the derivatives, namely rate of change and the others, do indeed always highlight and reveal time points of great interest. Unfortunately, while the time point is clear, it is very unclear as to which stations/buses are much more worthy of further and deeper attention and care and maybe intervention. For each of the scenarios provided, we have thoroughly explored the three measurements (voltage, frequency and phase angle) from two main perspectives, namely

1. **Novelty/Anomaly Detection:** This perspective of our statistical analysis views the $S = 126$ stations/buses provided as a random sample from a large population of stations. According to this view, at each time point t , we consider an S dimensional vector $(X_{1,t}, X_{2,t}, \dots, X_{s,t}, \dots, X_{S,t})$, but we also consider the $S \times t$ matrix of accumulated observations until time t .
2. **Time Series Clustering:** This view of the dynamics of the grid hypothesizes that there are groups of buses, probably determined by the internal level of instability. The assumption in this view is that there might be potentially two to three or even potential four types of buses, maybe (a) buses needing no intervention (stable) (b) buses starting to deteriorate (c) buses mildly out of order (d) buses severely out of order. At any rate, unlike novelty detection, this unsupervised classification (clustering) approach potentially wants to determine (detect and identify) groups of buses via appropriately designed statistical clustering algorithms.

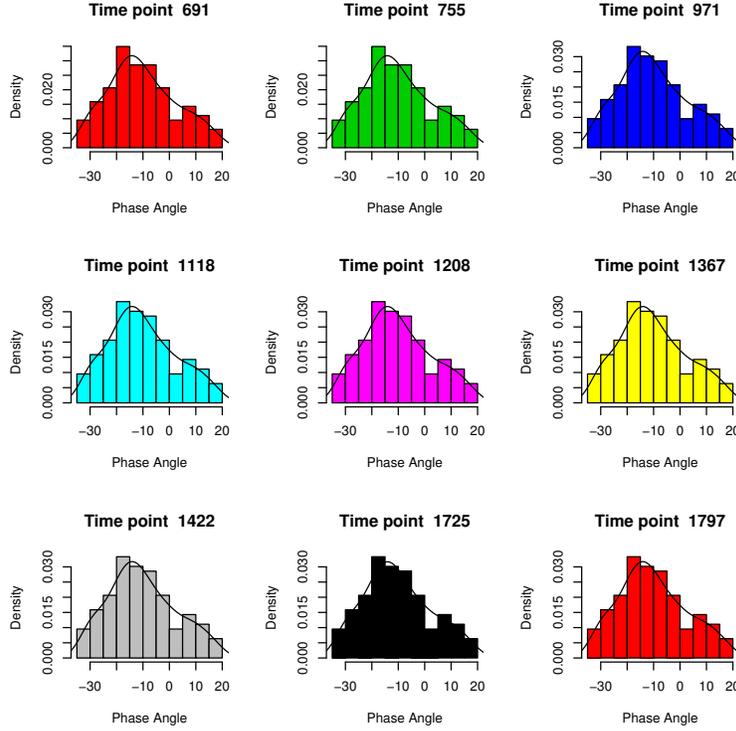


Figure 5. Measurements for the McNary attack scenario, featuring the raw series on phase angle across stations at fixed time points.

3.1 Computation of the elements of outlier detection

Our immediate solution to that problem came in the form of focusing (zeroing in) of time points of interest, which led to our first view of the BPA data matrix for each scenario, namely to analyze all the S series together at any given time point t . Therefore, we consider

$$X_t = (X_{1,t}, X_{2,t}, \dots, X_{s,t}, \dots, X_{S,t})^\top,$$

which is clearly the instantaneous consideration of all the S stations together at the single time point t .

At the risk of abuse of assumption, one could assume that X_t is a random sample of S observations drawn from the generating process at time t . This has the potential to help determine at each time point t , those observations, ie stations in this case, that might be exhibiting anomalies (novelties) worth addressing, thereby tackling and solving the problem of lack of station identification encountered earlier. It is interesting to start by look at samples of time points t to gain insights into how the measurements are distributed across buses. It is clear from Figure 5, that the instantaneous distribution of *phase angle* is bimodal and definitely not Gaussian. In fact, we even performed tests of normality that reveal the non-gaussianity for the McNary scenario.

In this case, we clearly see a mixture of two Gaussians. One the Gaussians obviously consistently has a larger proportion of observations along with a larger variance. This might refer to the regular (normal) regime, while the smaller Gaussian on the right with larger measurements might refer to potential anomalies.

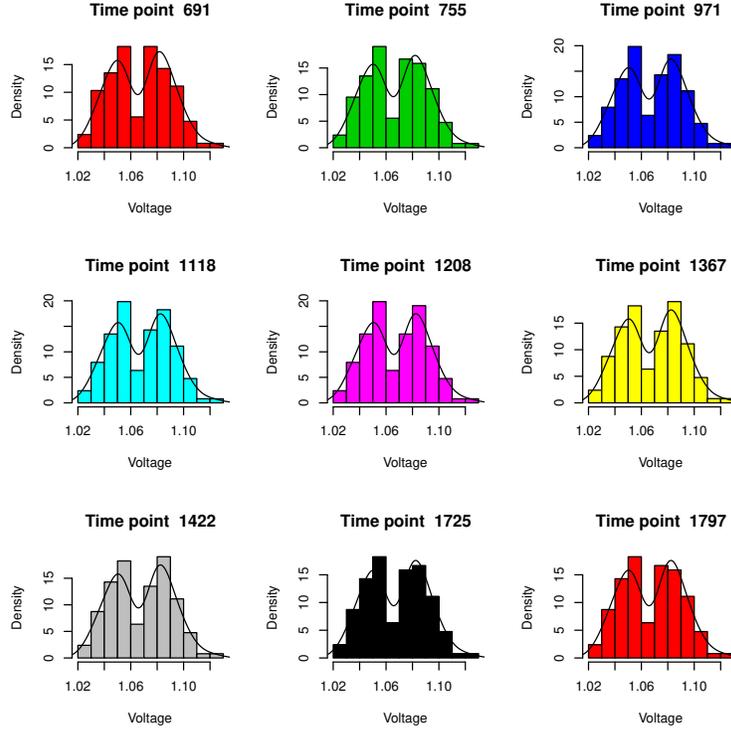


Figure 6. Measurements for the McNary attack scenario, featuring the raw series on voltage across stations at fixed time points.

Figure 6, reveals that the distribution of voltage is also bimodal, but unlike with phase angle, the two components of the mixture of Gaussians appear to be of equal probability. It is interesting to investigate the scientific reason behind the distributional differences. This be just an aspect of Voltage worth investigating further for the detection of potential anomalies.

From Figure 7, it appears that the distribution of frequency is quite unlike the previous two. It oscillates between uniform and sharply peaked at extremes. All the distributional considerations must be taken into account later to extract the highest quality of information from the collection of time series.

Now at any given time point t , we must consider the following

1. The instantaneous S dimensional vector $X_t = (X_{1,t}, X_{2,t}, \dots, X_{s,t}, \dots, X_{S,t})^\top$
2. The accumulated $S \times t$ matrix of all the t vectors of dimension S collected from time 1, namely

$$\mathbf{X}_{1:t} = \begin{bmatrix} X_{1,1} & X_{1,2} & \cdots & X_{1,t} \\ X_{2,1} & X_{2,2} & \cdots & X_{2,t} \\ \vdots & \ddots & \cdots & \ddots \\ X_{s,1} & X_{s,2} & \cdots & X_{s,t} \\ \vdots & \ddots & \cdots & \ddots \\ X_{S-1,1} & X_{S-1,2} & \cdots & X_{S-1,t} \\ X_{S,1} & X_{S,2} & \cdots & X_{S,t} \end{bmatrix} \in \mathbb{R}_+^{S \times t}$$

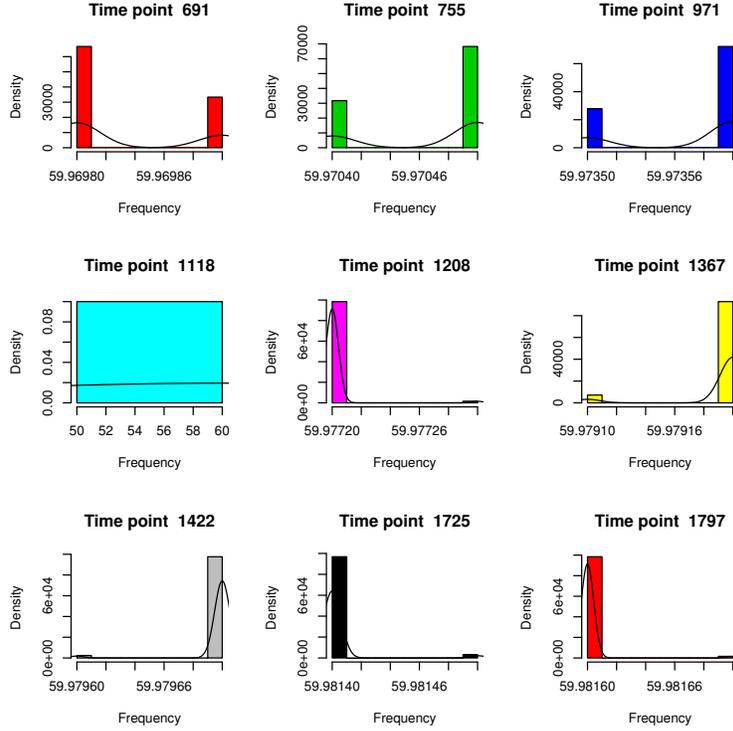


Figure 7. Measurements for the McNary attack scenario, featuring the raw series on frequency across stations at fixed time points.

3. The instantaneous S dimension vector $X_{t+1} = (X_{1,t+1}, X_{2,t+1}, \dots, X_{s,t+1}, \dots, X_{S,t+1})^\top$

$\mathbf{X}_{1:t}$ represents the accumulated measurements from beginning to time t in the form of a matrix for all the stations. Using that the sequence $\mathbf{X}_{1:t}, \mathbf{X}_{1:t+1}, \dots, \mathbf{X}_{1:T-1}, \mathbf{X}_{1:T}$ is just a dynamic monitoring of all the buses in time t .

Both Phase Angle and Voltage appear to be governed by a mixture of two Gaussians, judging from the marginal distributions plotted earlier. One could assume that the portion (component) with the lower mean is regular while the other component, the one with the higher mean is the potential source of anomaly. With that we can think of this as an ϵ -contamination regime, with the density of the observed measurement captured by

$$p(\mathbf{x}_t | \boldsymbol{\mu}_t, \Sigma_t, \epsilon, \eta, \gamma) = (1 - \epsilon)\phi_p(\mathbf{x}_t; \boldsymbol{\mu}_t, \Sigma) + \epsilon\phi_p(\mathbf{x}_t; \boldsymbol{\mu}_+ \eta, \gamma \Sigma) \quad (3.1)$$

where η and γ represent the level of contamination of the location and scatter matrix, respectively, and

$$\phi_p(\mathbf{x}_t; \boldsymbol{\mu}_t, \Sigma_t) = \frac{1}{\sqrt{(2\pi)^p |\Sigma_t|}} \exp \left\{ -\frac{1}{2} (\mathbf{x}_t - \boldsymbol{\mu}_t)^\top \Sigma_t^{-1} (\mathbf{x}_t - \boldsymbol{\mu}_t) \right\}. \quad (3.2)$$

Although occasional drifts were observed especially with Frequency measurements, none of the series ever shows any trend at all, which is to be expected, given that the grid is in good working order almost all the time, and the problems being monitoring are rare events triggering extreme measurement values. With no trend to contend with and clearly no seasonal component in the traditional sense of seasonality, we are left

with a random component being observed over time. The most important questions are: (a) Are we in the presence of stationary time series? (b) Do these series lend themselves to modelling with the ARMA family of stochastic processes? If so, which members of the family are optimal?

Essentially, $\boldsymbol{\mu}_t = \mathbb{E}[X_t]$ and $\Sigma_t = \text{cov}(X_t)$, which we estimate based the accumulated data $\mathbf{X}_{1:t}$ using

$$\widehat{\Sigma}_t = \frac{1}{t-1} \sum_{i=1}^t (X_i - \widehat{\boldsymbol{\mu}}_t)(X_i - \widehat{\boldsymbol{\mu}}_t)^\top$$

where $X_i = (X_{1,i}, X_{2,i}, \dots, X_{s,i}, \dots, X_{S,i})^\top$ and $\widehat{\boldsymbol{\mu}}_t = \frac{1}{t} \sum_{i=1}^t X_i$. In the interest of computational efficiency and stability, we dynamically (sequentially) update the estimates of the mean and the scatter matrix at time $t+1$ using

$$\widehat{\Sigma}_{t+1} = \frac{t-1}{t} \widehat{\Sigma}_t + \frac{1}{t+1} \mathbf{x}_{t+1} \mathbf{x}_{t+1}^\top.$$

Note that under stationarity, both the mean and the scatter matrix will NOT depend on time, and the updating will simply be an improvement of the time independent $\boldsymbol{\mu}$ and Σ . Now, for outlier detection, since each X_t is S -dimension, we have

$$(X_{t+1} - \boldsymbol{\mu}_t)^\top \Sigma_t^{-1} (X_{t+1} - \boldsymbol{\mu}_t) \sim \chi^2(S)$$

which is a chi-squared distribution with S degrees of freedom. We can therefore declare

1. The vector X_{t+1} seriously (extremely) problematic and needing intervention if

$$(X_{t+1} - \widehat{\boldsymbol{\mu}}_t)^\top \widehat{\Sigma}_t^{-1} (X_{t+1} - \widehat{\boldsymbol{\mu}}_t) > \chi_{0.01}^2(S)$$

2. The vector X_{t+1} somewhat problematic and needing intervention if

$$(X_{t+1} - \widehat{\boldsymbol{\mu}}_t)^\top \widehat{\Sigma}_t^{-1} (X_{t+1} - \widehat{\boldsymbol{\mu}}_t) > \chi_{0.05}^2(S)$$

3. The vector X_{t+1} regular and not needing any intervention if

$$(X_{t+1} - \widehat{\boldsymbol{\mu}}_t)^\top \widehat{\Sigma}_t^{-1} (X_{t+1} - \widehat{\boldsymbol{\mu}}_t) < \chi_{0.05}^2(S)$$

- (a) We have estimated at each time point t the distribution of the measurements (voltage, frequency, phase angle) and have so far at least informally established that *voltage* and *phase angle* are both mixtures of two Gaussians, while *frequency* appears to exhibit a point mass distribution (or at times uniform) making it far clearer to separate the wheat from the chaff.
- (b) We have been able with both mixtures of Gaussians and nonparametric density estimators along with available univariate outlier detectors to detect stations needing intervention
- (c) For each time point t , we have detected outliers using simple and computational efficient criteria like

$$\text{if } X_{st} \notin [\mathbf{Q}_1 - 3 \times \text{IQR}, \mathbf{Q}_3 + 3 \times \text{IQR}]$$

then declare station s a severe outlier needing intervention. The identity of stations along with the magnitude of the deviation are used in the construction of the heatmap displayed to the operator. Notice that we did NOT use the univariate Gaussian outlier detection that declares station s an outlier

$$\text{if } X_{st} \notin [\widehat{\boldsymbol{\mu}}_t - 3 \times \sqrt{\widehat{\sigma}_t^2}, \widehat{\boldsymbol{\mu}}_t + 3 \times \sqrt{\widehat{\sigma}_t^2}]$$

Our reason for not using the Gaussian criterion is due to the fact the distributions revealed by the histograms were mixtures of Gaussians and not Gaussians.

- (d) Realizing that the time points are indeed related since we are in the presence of a time series data with inherent serial temporal correlations, we also crucially considered and explored the accumulated data until time t to learn from the data prior to detecting the outliers. The accumulated $S \times t$ matrix of all the t vectors of dimension S collected from time 1, namely

$$\mathbf{X}_{1:t} = \begin{bmatrix} X_{1,1} & X_{1,2} & \cdots & X_{1,t} \\ X_{2,1} & X_{2,2} & \cdots & X_{2,t} \\ \vdots & \ddots & \cdots & \ddots \\ X_{s,1} & X_{s,2} & \cdots & X_{s,t} \\ \vdots & \ddots & \cdots & \ddots \\ X_{S-1,1} & X_{S-1,2} & \cdots & X_{S-1,t} \\ X_{S,1} & X_{S,2} & \cdots & X_{S,t} \end{bmatrix} \in \mathbb{R}_+^{S \times t}$$

Instead of looking at the time point t in isolation as we did so far, we determine the quality of the time point t in relation to the accumulated measurements until time t . For outlier detection under the multivariate Gaussian distribution, since each X_t is S -dimension, we have

$$(X_{t+1} - \boldsymbol{\mu}_t)^\top \Sigma_t^{-1} (X_{t+1} - \boldsymbol{\mu}_t) \sim \chi^2(S)$$

which is a Chi-squared distribution with S degrees of freedom.

We can therefore do the following:

1. Declare the vector X_{t+1} seriously (extremely) problematic and needing intervention if

$$(X_{t+1} - \hat{\boldsymbol{\mu}}_t)^\top \hat{\Sigma}_t^{-1} (X_{t+1} - \hat{\boldsymbol{\mu}}_t) > \chi_{0.005}^2(S)$$

2. Declare the vector X_{t+1} somewhat problematic and needing intervention if

$$(X_{t+1} - \hat{\boldsymbol{\mu}}_t)^\top \hat{\Sigma}_t^{-1} (X_{t+1} - \hat{\boldsymbol{\mu}}_t) > \chi_{0.025}^2(S)$$

3. Declare the vector X_{t+1} regular and not needing any intervention if

$$(X_{t+1} - \hat{\boldsymbol{\mu}}_t)^\top \hat{\Sigma}_t^{-1} (X_{t+1} - \hat{\boldsymbol{\mu}}_t) < \chi_{0.025}^2(S)$$

where under the multivariate Gaussian assumption,

$$\hat{\Sigma}_t = \frac{1}{t-1} \sum_{i=1}^t (X_i - \hat{\boldsymbol{\mu}}_t)(X_i - \hat{\boldsymbol{\mu}}_t)^\top$$

where $X_i = (X_{1,i}, X_{2,i}, \dots, X_{s,i}, \dots, X_{S,i})^\top$ and $\hat{\boldsymbol{\mu}}_t = \frac{1}{t} \sum_{i=1}^t X_i$.

$$\hat{\Sigma}_{t+1} = \frac{t-1}{t} \hat{\Sigma}_t + \frac{1}{t+1} \mathbf{x}_{t+1} \mathbf{x}_{t+1}^\top.$$

For the purpose of generating a compelling visualization revealing all the anomalous stations/buses to the operator, all the stations are colored according to their contribution to the Mahalanobis distance $d_{\Sigma_t}(X_{t+1}, \boldsymbol{\mu}_t)$ of X_{t+1} to the mean $\boldsymbol{\mu}_t$ of the process, where

$$d_{\Sigma_t}(X_{t+1}, \boldsymbol{\mu}_t) = (X_{t+1} - \hat{\boldsymbol{\mu}}_t)^\top \hat{\Sigma}_t^{-1} (X_{t+1} - \hat{\boldsymbol{\mu}}_t).$$

Once $X_{t+1} = (X_{1,t+1}, X_{2,t+1}, \dots, X_{s,t+1}, \dots, X_{S,t+1})^\top$ is declared anomalous, The level of severity is obtained from the ranges given earlier.

1. The components of the vector $\widehat{\Sigma}_t^{-1}(X_{t+1} - \widehat{\boldsymbol{\mu}}_t)$, which are buses, are extracted
2. The components of the vector $\widehat{\Sigma}_t^{-1}(X_{t+1} - \widehat{\boldsymbol{\mu}}_t)$, which are buses, are scored
3. The score for each station is then mapped to the color of that bus
4. The coordinate of the buses are extracted
5. Each of the station is re-mapped using computationally intensive polygonal calculations
6. The color is then assigned on the map to each station based on their score. Extremely deviating stations receive levels of red, mildly deviating stations receive levels of yellow, and normal (non deviating) buses receive levels of green.

3.2 Visualization of Anomalies on Heat Maps

The key to this approach is a mix of univariate nonparametric novelty detection and parametric outlier detection. The heat map generation does not look at the substations, but instead focuses entirely on each bus as its own entity. Essentially, the visualization procedure is as follows:

1. The geographic coordinates of each bus are extracted along with its contours
2. A polygon is constructed to represent the bus region on the map
3. The values of the measurements are matched with colors from the set [red, yellow, green, blue, black] with red representing the lowest value and black representing the largest values, and the matching is made richer and better revealing by the use of gradient filling, which makes the coloring on the map continuous
4. Rate of change turns out to be the most revealing of the function of the measurements

The creation of the heat maps is computationally very heavy (expensive/cost) because of the need to generate/create the polygons that represent the region of the bus on the map. The following is the R code we used to create the heat maps

```
#
# Plot the heatmap of the rate of change of the phase angle, at time <t>, and
# returns the ggplot object representing that plot
#

plot_heatmapangle_firstroc <- function(t)
{
  tf <- t(t(firstroc[t,-1]))
  tf <- cbind(rownames(tf),tf)
  colnames(tf) <- c("Bus.Name", "Angle")
  bus_loc_angle_foroc <- merge(bus_locs[ , ! colnames(bus_locs) %in% c("Angle")],tf, by="Bus.Name")
  bus_locs$Angle <- as.numeric(as.character(bus_loc_angle_foroc$Angle))
  # bus_locs <- update_freq_firstroc(t)

  if(!exists("autosc")){
    autosc <-<- FALSE
  }
}
```

```

}
if(!exists("is_zoom")){
  is_zoom <- 0
}
if (is_zoom!=0) {
  xmn <- min(bus_locs[(bus_locs$Longitude>=z_xmin),"Longitude"])
  xmx <- max(bus_locs[(bus_locs$Longitude<=z_xmax),"Longitude"])
  ymn <- min(bus_locs[(bus_locs$Latitude>z_ymin),"Latitude"])
  ymx <- max(bus_locs[(bus_locs$Latitude<z_ymax),"Latitude"])
} else{
  xmn <- min(bus_locs$Longitude)
  xmx <- max(bus_locs$Longitude)
  ymn <- min(bus_locs$Latitude)
  ymx <- max(bus_locs$Latitude)
}
xstep <- (xmx-xmn)/80
ystep <- (ymx-ymn)/80

intp_coords <- interp(bus_locs$Longitude, bus_locs$Latitude, bus_locs$Angle, duplicate = "mean",
                    xo=seq(xmn,xmx, by=xstep),
                    yo=seq(ymn,ymx, by=ystep))
r <- raster(intp_coords)
rtp <- rasterToPolygons(r)

rtp@data$id <- 1:nrow(rtp@data) # add id column for join
rtpFort <- fortify(rtp, data = rtp@data)
rtpFortMer <- merge(rtpFort, rtp@data, by.x = 'id', by.y = 'id') # join data

amin <- ifelse(min(bus_locs$Angle)<(-0.001),min(bus_locs$Angle),(-0.001))
amax <- ifelse(max(bus_locs$Angle)>0.001,max(bus_locs$Angle),0.001)
g <- g + geom_polygon(data = rtpFortMer,
                    aes(x = long, y = lat, group = group, fill = layer),
                    alpha = 1,
                    size = 0,
                    color = NA) + ## size = 0 to remove the polygon outlines
scale_fill_gradientn("Rate of Change",colours = c("red","yellow","green","blue","black"),limits=c
theme(legend.position="right",legend.direction="vertical",legend.box="horizontal",axis.title.x=e
      axis.text.x=element_blank(),
      axis.ticks.x=element_blank(),
      axis.title.y=element_blank(),
      axis.text.y=element_blank(),
      axis.ticks.y=element_blank())
# ggtitle(bquote(atop("Phase Angle at Time",atop(. (Pangle[t,1]),""))))
g
}

```

Figure 8 shows an example of a heat map visualization of rate of change of phase angles in the open AC line scenario. It is easy to spot the anomaly in southwest Oregon by the different color, which is the designed

advantage of geographically-based color-contoured maps. We chose to display rate of change for the same reason as in the previous visualization: Rate of change “filters” normal variability between buses highlighting *change*, which should be of interest to the operators, and possibly allows for anticipation of the future state of the system, supporting Level 3 SA.

Our approach in the generation of the heat map visualizations was purely statistical and heavily data analytic, in the sense that we did not rely on any assumptions about the physics of the measurements under consideration. For example, to gain insights into the pattern underlying each measurement we first considered the raw measurement itself and then looked at score of variation/fluctuation such as rate of change, first order difference and second order difference. It is known in time series analysis that differencing helps reveal the pure stochastic component from a series by removing components like trend, seasonality, and cycle.

Note, too, that Figures 8, 9, and 10 are still images of dynamic displays, depicting the measured values only in one instant of time. As the data comes in in operational setting the colors change in time. Such movement on the display has been shown to be very powerful in attracting attention to it, which in this case means a potential anomaly in system demanding operators’ attention.

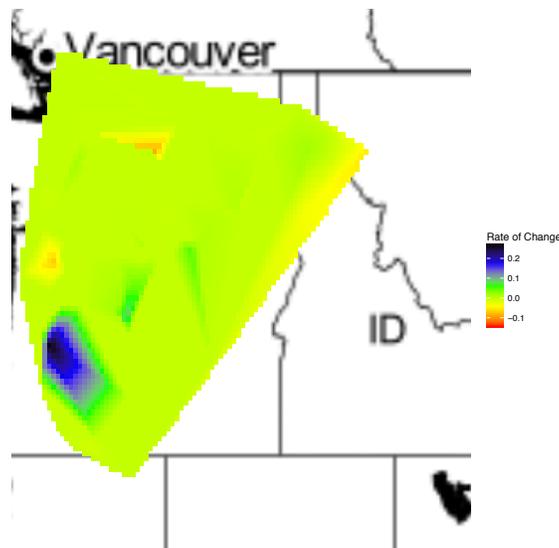


Figure 8. An example of a heat map visualization of phase angles in the open AC line scenario.

Figures 9 and 10 show the heat map visualization for all the key variables (phase angles, frequency, and voltage) in all the simulated scenarios. Although it is hard to tell from the still, instantaneous, images, we expect different grid events to create distinct dynamic color patterns that experienced operators would not only immediately notice but also identify, hastening their taking appropriate actions in response to the event.

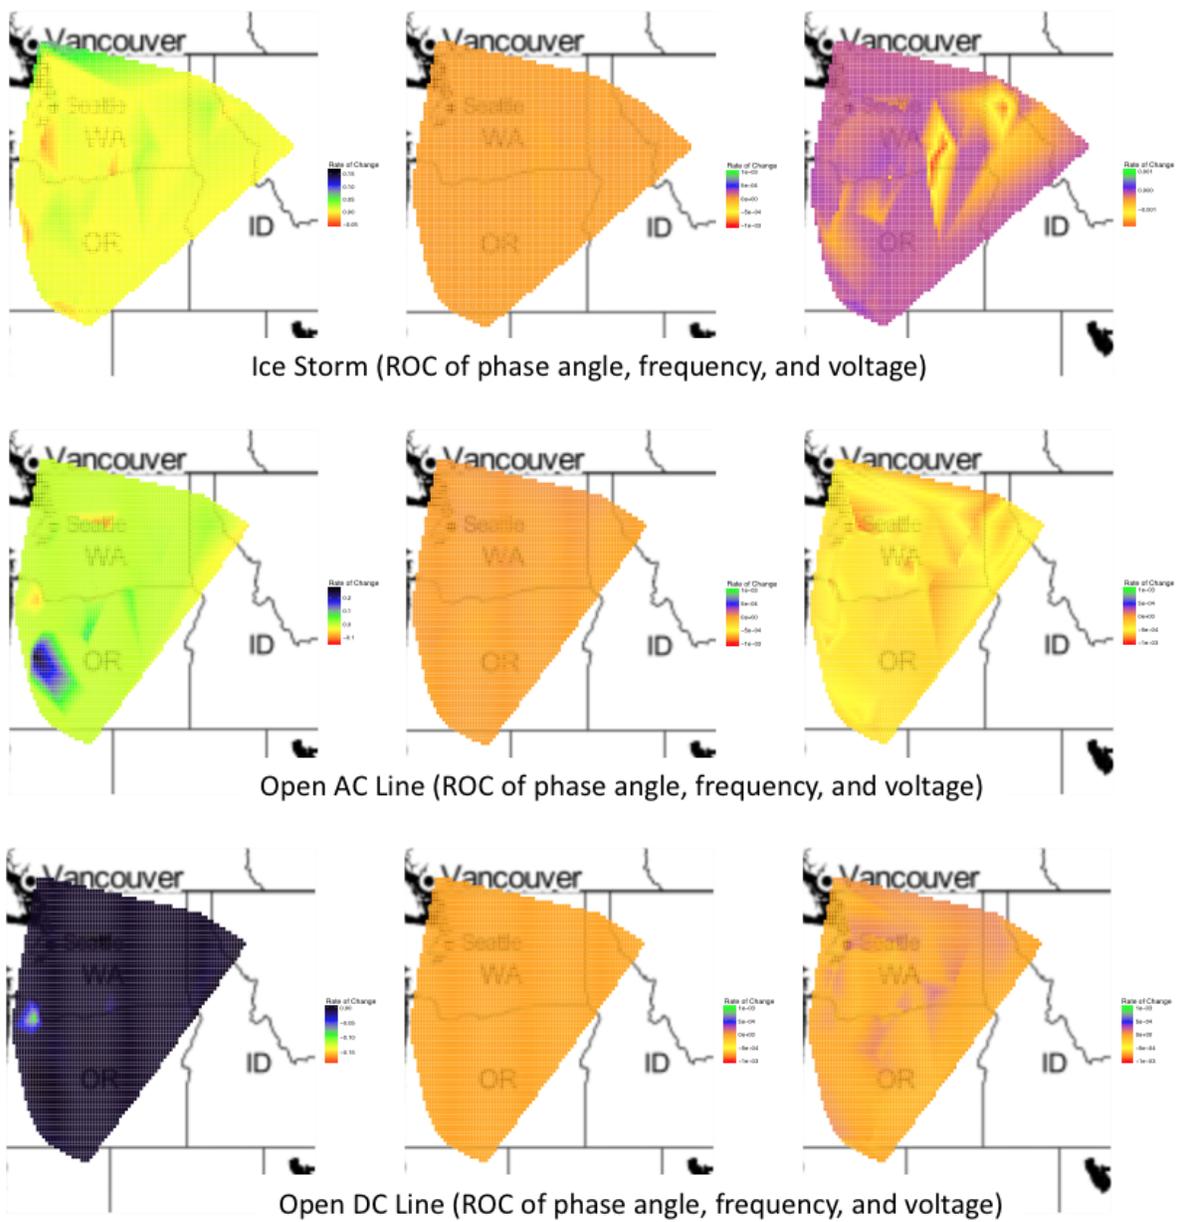


Figure 9. Example of a heat map visualizations all the key variables (phase angles, frequency, and voltage) in three of the simulated scenarios.

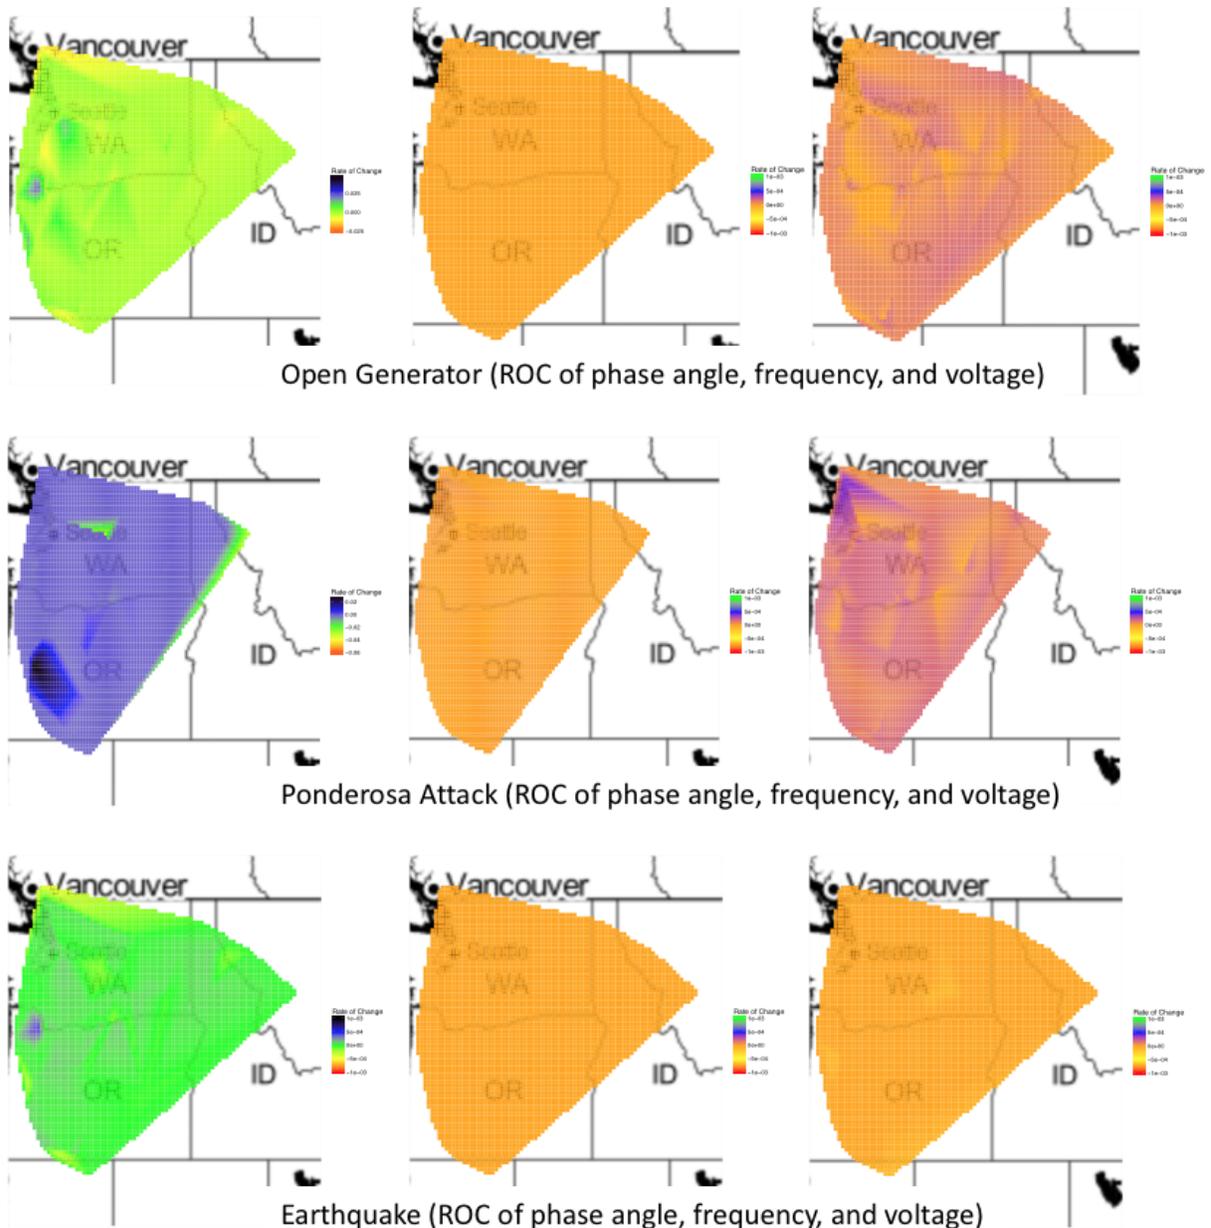


Figure 10. Example of a heat map visualizations all the key variables (phase angles, frequency, and voltage) in the other three simulated scenarios.

3.3 Time Series Clustering of Stations/Buses

As suggested by some of the histogram plotted earlier, one could rightly hypothesize that buses and substations can be clustered in two or three main groups, plausibly representing different levels of states of the electric grid. In this section, we explore the preliminary clustering of buses using distances and similarity measures that are specifically designed for the kind of time series yielded by the data from the buses. We specifically used various complexity-invariant similarity measures to do the following:

1. Generate the clustering/grouping of the buses into meaningful or revealing categories at each time

point t or a given window of time

2. Discover or extract or learn aspects of the buses from the different scenarios provided thus far. We think of the second part as classification in pattern recognition sense of the different scenarios with the ultimately goal of hopefully identifying a small number of types of situations to expect from the buses.

We considered the kMeans clustering algorithm and the agglomerative hierarchical clustering with various distances, and also explored an extension of the kMeans clustering algorithm known as the kMedoids clustering algorithm. In this case, we specifically used the function `pam` from the package `cluster`. `pam` stands for partitioning around medoids, which really replaces the mean used in kMeans with a single observation from \mathcal{D} and then replaces the ℓ_2 distance in kMeans, with any distance deemed appropriate for the data at hand.

Let $\mathcal{P}_k = \{(C_j, \mathbf{y}_j), j = 1, \dots, k\}$ represent the partition of the dataset into k clusters C_1, C_2, \dots, C_k with medoids $\{\mathbf{y}_1, \dots, \mathbf{y}_k\} \subset \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, the partitioning around k Medoids seek to minimize the within cluster dissimilarities with respect to some distance $d(\cdot, \cdot)$, namely

$$\mathcal{P}_k^* = \arg \min_{\mathcal{P}_k} \left\{ \text{WCD}(\mathcal{P}_k) \right\} = \arg \min_{\mathcal{P}_k} \left\{ \sum_{j=1}^k \sum_{\mathbf{x}_i \in C_j} d(\mathbf{x}_i, \mathbf{y}_j) \right\}$$

If the ℓ_2 distance is used, and the medoids are collapsed to the center, we have kMeans clustering. Clearly the kMeans clustering method cannot be used right off the shelf for the kind of time series data contained from the buses, because with all the huge variability and a great potential for outliers, along with the series having different lengths, the generic kMeans clustering algorithm is not appropriate. Besides, the use of Kmeans presupposes that the Euclidean distance is used on these time series without any adaptation, which is not appropriate. The clustering of these time series is non-trivial. Among other things, we need to emphasize:

1. Explore existing and potentially create the most appropriate distances and/or similarity measures that are best suited to the type of time series inherent in the PMU data
2. Implement brand new or adapt existing clustering techniques to achieve within series clustering and between series clustering with the ultimate goal of spotting the most striking patterns to be revealed to the operator for action (This may be algorithms like Partitioning Around Medoids using the most suitable distance measures, or one could even use nonparametric approaches to clustering where the number of clusters does NOT have to be specified
3. Revisit and properly report on the direct novelty detection in the dynamic progress of the buses/stations
4. And since we appear to be seriously considering the interactions between buses, we should perform some network analysis, even a basis one to potentially detect communities of buses dynamically. This may appear like typically clustering, but it is actually somewhat different and exploit beneficial aspects of graph-theoretic analysis.

Many authors have explored and applied time series distances in various settings. For example, [2] performed a *comparison of similarity measures for trajectory clustering in outdoor surveillance scenes*. As far as our current collection of BPA data scenarios is concerned, we chose the McNary attack scenario as our example to explore the strengths and weaknesses of some of the most commonly used dissimilarity measures for time series. Among all the distances consider in our implementation, we herein give more details on four of them:

1. **The Euclidean distance:** Given two time series $X = (x_1, x_2, \dots, x_p)$ and $Y = (y_1, y_2, \dots, y_q)$, the Euclidean distance between them is given by

$$d_E(X, Y) = \sqrt{\sum_{j=1}^{\min(p,q)} |x_j - y_j|^2}.$$

As can be seen, the Euclidean distance requires the two series to be of the same length, hence our use of $\min(p, q)$ to guarantee that the portions participating in the computation are of the same length. We thought it useful to mention the Euclidean distance here despite its inherent inadequacy for time series, and we mentioned it primarily as a baseline approach to be used for comparison purposes.

2. **The Fréchet distance:** A recent treatment of Time Series Clustering using the Fréchet distance can be found in [3]. The following site <http://dtw.r-forge.r-project.org/>, provides a very detailed implementation. To better understand it, let us now give a working definition of the Fréchet distance. Consider two time series $X = (x_1, x_2, \dots, x_p)$ and $Y = (y_1, y_2, \dots, y_q)$. Let r be a sequence of m pairs preserving the observation order, such that

$$r = ((x_{a_1}, y_{b_1}), (x_{a_2}, y_{b_2}), \dots, (x_{a_m}, y_{b_m}))$$

where $a_i \in \{1, 2, \dots, p\}$, $b_j \in \{1, 2, \dots, q\}$ with $a_1 = 1, a_m = p, b_1 = 1, b_m = q$ and $a_{i+1} = a_i$ or $a_{i+1} = a_i + 1, b_{i+1} = b_i$ or $b_{i+1} = b_i + 1$ for $i \in 1, 2, \dots, m - 1$. The Fréchet distance $d_F(X, Y)$ between X and Y is defined as

$$d_F(X, Y) = \min_{r \in M} |r| = \min_{r \in M} \left\{ \max_{i \in [m]} |x_{a_i} - y_{b_i}| \right\} \quad (3.3)$$

3. **Longest Common Subsequence (LCSS):** The distance is essentially defined recursively (recurrently) through an algorithm. Essentially, let $X = (x_1, x_2, \dots, x_p)$ and $Y = (y_1, y_2, \dots, y_q)$. Then let $i = 1, \dots, p$ and $j = 1, \dots, q$. Then,

$$\text{LCSS}(X, Y) = \begin{cases} 0 & \text{if } |X| = 0 \text{ or } |Y| = 0 \\ 1 + \text{LCSS}(X_{i-1}, Y_{j-1}) & \text{if } \text{distance}(X_i, Y_j) < \gamma \\ \max(\text{LCSS}(X_{i-1}, Y_j), \text{LCSS}(X_i, Y_{j-1})) & \text{Otherwise.} \end{cases}$$

4. **The Dynamic Time Warping (DTW):** This distance is also defined recursively via an algorithm. Let us once again consider two time series vectors $X = (x_1, x_2, \dots, x_p)$ and $Y = (y_1, y_2, \dots, y_q)$, and let $i = 1, \dots, p$ and $j = 1, \dots, q$. An essential component of the distance $\text{DTW}(X, Y)$ is $(D(i, j))$, the Dynamic Time Warping distance between the partial/sub sequences $X_{1:i}$ and $Y_{1:j}$, and is given by

$$D(i, j) = |X_i - Y_j| + \min \begin{bmatrix} D(i-1, j) \\ D(i-1, j-1) \\ D(i, j-1) \end{bmatrix}$$

As we can see above, most of the distance are computationally quite complex, somewhat reflecting the fact that time series are not simple vectors in the traditional sense, but rather complicated constructs that should be handled with accordingly complex and appropriate measures of similarities. Now, for our implementation of the time series distances, we fittingly resorted to various R packages dedicated to either similarity measures for time series or comprehensive time series clustering. The R package `SimilarityMeasures` for instance has functions for the following time series distances: Longest Common Subsequence (LCSS), Dynamic

Time Warping (DTW), Fréchet Distance, and the Edit Distance.

The R package `TSdist` is even more comprehensive, providing many more similarity measures than `SimilarityMeasures`. The most comprehensive R package to date for clustering time series data is `TSclust` presented in extensive detail by [4]. The most complete reviews of all the state-of-the-art techniques and methods of time series clustering can be found in [5] and [6]. Another package along the same lines is explored in [7]. In the neural networks literature [8] have developed and implemented various techniques of time series clustering. Deep Learning is quickly becoming the default approach for solving pretty much every simple problem in statistical machine learning and artificial intelligence. Some authors have developed adaptations of deep learning paradigm for time series (e.g., [9]).

Finally, for completeness sake, it is worth mentioning that the setting of time series clustering could also be appropriately formulated as an anomaly detection problem, for which a flourishing literature has emerged recently as capturing in [10]. The subsequence time series clustering approach found in [11] has also proven to be successful in some settings.

We performed the clustering of (1) voltage (2) frequency, and (3) phase angle, all of those on the raw data. The Fréchet distance and other distances mentioned earlier are not shown in this report, because those distances are computationally very intensive, requiring days of uninterrupted computing for large/long series like ours. We did explore the performances of those distances on our data, but such computationally intensive approaches are not feasible for operational implementation of the visualizations, where streaming data must appear on displays with minimum lags.

It is worth emphasizing that we concentrated for now solely on hierarchical clustering, mainly because of the visual appeal and interpretability of the dendrogram that results from hierarchical clustering. Given that we have a wide variety of distances available to us, we plan to also consider implementing k-Medoids approach, namely Partitioning Around Medoids (PAM).

It appears from the dendrograms we examined that frequency is the measure that seems to exhibit a clear separation of the buses into groups. One can see with all the 5 distance used on frequency, that there appears to be a clear separation into two groups. It was crucial to process the labels of the groups a posteriori to try to identify the buses falling into each of the two groups:

1. Examine (carefully) the characteristics of the members of each clusters, with the hope of finding the problematic stations in one group and the steady-state stations in the other.
2. Perform the sequential clustering of frequency to find out if what we saw with the whole data can be tracked sequentially.
3. Consider another scenarios, and find out if frequency tends to be the measure for which separations are more clear cut.
4. Generate experimental scenarios based on that revealing feature of frequency.

It can also be seen, that with the voltage and phase angle, the separations of the clusters is not as clear cut as with frequency.

1. Each station is treated as its own time series, which makes the use of histograms an inadequate graphical summary. What we have done so far is the exploration of (a) Spectral density of the time series of each station (b) the autocorrelation function (ACF) (c) the partial autocorrelation function (PACF).

Voltage and Phase Angle reveal the potential for ARIMA modeling. Frequency however is more complex.

2. We have constructed the clustering of stations using a wide variety of time series specific distances and it appears that sharper separations are gained for frequency far more than voltage and phase angle.

The key to this approach in clustering followed by univariate nonparametric outlier detection. At a given time t ,

1. Perform cluster analysis on the measurements, specifically extracting a 3-clusters solution as suggested by most of the histograms
2. Obtain the average/center/mean of each of the three clusters
3. Rank order the absolute values of the 3 centers in increasing order
4. Focus on the cluster with the largest absolute center, since we are seeking to identify those buses and the substations thereof that have worrisome fluctuations worthy (needing to be) of being addressed by the operator
5. If the sign of the center is negative, then we are in the presence of what we call down measurement, otherwise we have an up measurement
6. Within the chosen cluster, consider all the buses therein
 - (a) Extract all buses such that their measurements are outside of the interval $[Q1-1.5 IQR, Q3+1.5 IQR]$ for mild deviations (fluctuations) or outside of $[Q1-3 IQR, Q3+3 IQR]$ for extreme deviations
 - (b) For all the buses outside the normal as described above, identify their substation
 - (c) For substations with multiple buses, choose the representative as follows: if we are in down measurement regime, pick the smallest value, and if we are in up measurement regime, pick the largest value
 - (d) Identify the latitude and longitude of the substation and place the pin there. (different colors and sizes depending on the measurement (Phase Angle, Frequency, Voltage))

The creation of pinned maps is computationally very efficient because it suffices for us to consider the univariate time series made up of the B measurements for each of the B buses. The computational efficiency goes even further in the sense that, unlike heatmaps for which very heavy computations are needed to fill up polygons on the maps, the pinned maps is computationally “free of charge”, yet provides the information needed by the operator.

Figure 11 shows an example of a pinned map in the open generator scenario. Note that as was the case with the heat maps, time in this display is represented dynamically by animation (i.e., the pins appear and disappear in time as the grid event unfolds); the static Figure 11 only shows one instant in time. Nonetheless, we expect different events to produce different patterns of pins that may be recognizable by experienced operators.

The pins, too, could be formatted to convey more information than what the (default) pin symbols in Figure 11 do. For example, pin height could be used to code the magnitude of the anomaly at each bus, or the pins could have labels corresponding to the bus names, further directing the operators’ attention to the exact location of the problem. Great care must be taken that such additional information put on the display does not clutter the display and render it unreadable.

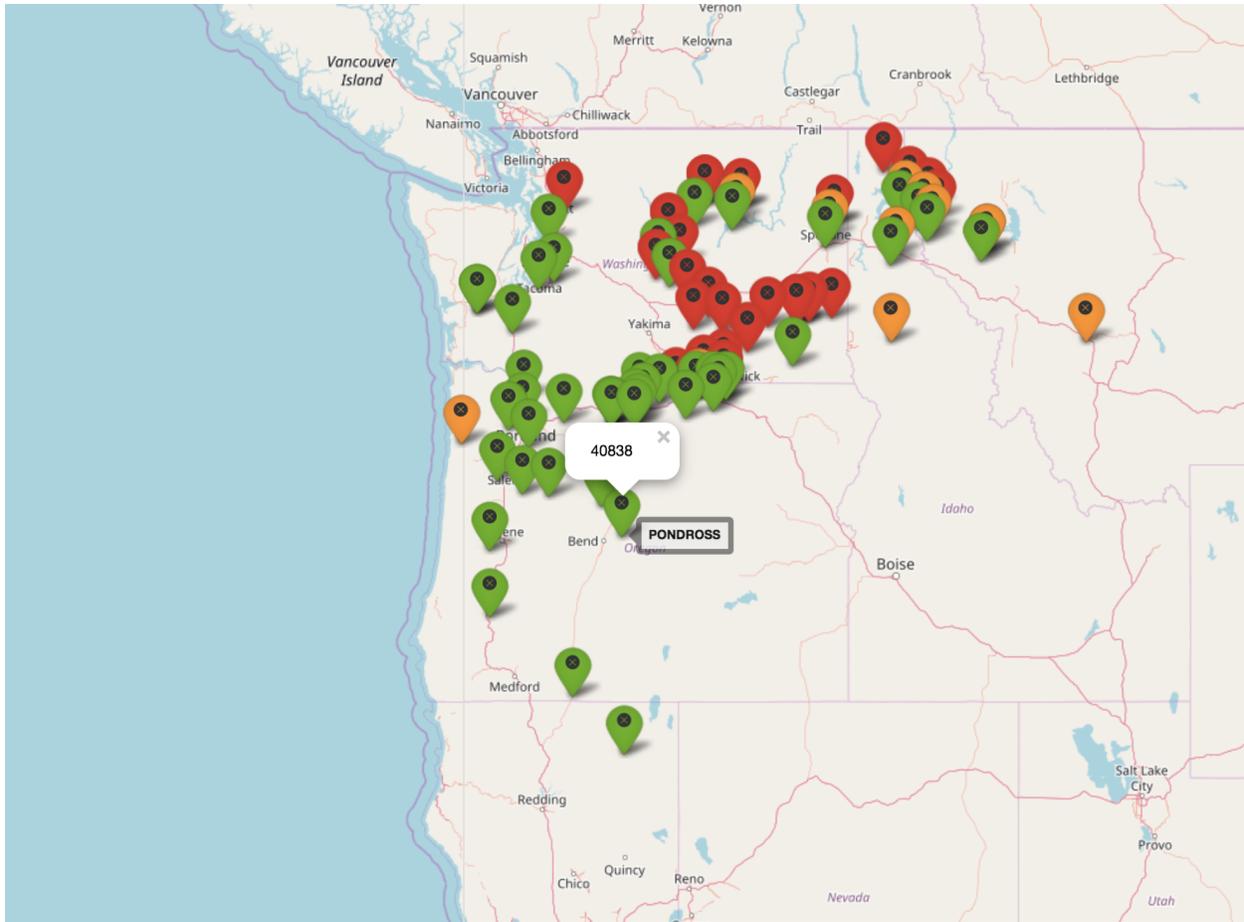


Figure 11. An example of a pinned map visualization of clusters of buses exhibiting anomalous behavior at a given time. Voltage, frequency, and phase angles are represented by the color of the pins.

References

- [1] L. Wilkinson. *The grammar of graphics*. Springer Science & Business Media, 2006.
- [2] Z. Zhang, K. Huang, and T. Tan. Comparison of similarity measures for trajectory clustering in outdoor surveillance scenes. In *Proceedings of the 18th International Conference on Pattern Recognition - Volume 03, ICPR '06*, pages 1135–1138, Washington, DC, USA, 2006. IEEE Computer Society.
- [3] A. Driemel, A. Krivosija, and C. Sohler. Clustering time series under the Fréchet distance. *CoRR*, abs/1512.04349, 2015.
- [4] P. Montero and J. A. Vilar. TSclust: An R package for time series clustering. *Journal of Statistical Software*, 62(1):1–43, 2014.
- [5] S. Aghabozorgi, S. A. Shirkhorshidi, and T. Ying Wah. Time-series clustering - a decade review. *Inf. Syst.*, 53(C):16–38, October 2015.
- [6] T. Warren Liao. Clustering of time series data-a survey. *Pattern Recognition*, 38(11):1857–1874, November 2005.

- [7] A. M. Brandmaier. pdc: An R Package for Complexity-Based Clustering of Time Series. *Journal of Statistical Software*, 67(5):1–23, 2015.
- [8] S. L. Suárez Gómez, J. D. Santos Rodríguez, F. J. Iglesias Rodríguez, and F. J. de Cos Juez. Analysis of the temporal structure evolution of physical systems with the self-organising tree algorithm (sota): Application for validating neural network systems on adaptive optics data before on-sky implementation. *Entropy*, 19(3), 2017.
- [9] J. C. B. Gamboa. Deep learning for time-series analysis. *CoRR*, abs/1701.01887, 2017.
- [10] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):15:1–15:58, July 2009.
- [11] K. A. Peker. Subsequence time series (STS) clustering techniques for meaningful pattern discovery. In *International Conference on Integration of Knowledge Intensive Multi-Agent Systems*, pages 360–365, April 18–21 2005.

Part IV

Visualization of Modes and Sources of Oscillations

1 Modal Analysis Background

Modal analysis is a well-known technique applied in power systems. It is used in assessing the small signal stability of the system (ability to withstand small disturbances), as well as to determine characteristics of electromechanical oscillations. In this part of the report, we present novel visualizations of power system modes and oscillations, to enable better situational awareness and decision making.

1.1 Methods

Broadly, modal analysis techniques can be categorized into two types –

1. Model based methods: This involves linearizing the entire power system model at an operating point. A detailed representation of the power grid is needed. In other words, the power grid model, along with the dynamic models for its generators and control devices such as machines, exciters, governors, stabilizers and load models have to be included.
2. Signal based methods: Modes can be determined from analyzing measurements from the grid, such as PMU data. This can be done totally independent of the model. Methods such as Prony Analysis [1], Matrix Pencil [2],[3], Variable Projection Method [4], etc. are examples of signal based, model-less methods. These are the focus of our work.

The basic idea of all signal based methods is to approximate a signal, $y_{\text{org}}(t)$, by the sum of other, simpler signals, which are also known as basis functions. Some properties of basis functions are -

- Basis functions are usually exponentials, with linear and quadratic functions also added to de-trend the signal.
- Properties of the original signal can be quantified from basis function properties (such as frequency and damping).
- Signal is considered over an interval with $t=0$ at the beginning.

The original signal $y_{\text{org}}(t)$ is sampled to generate a vector \mathbf{y} consisting of m uniformly sampled points. Let the sampling value be ΔT , starting with $t=0$, with values y_j for $j = 1 \dots m$. At each time point j , where $t_j = (j - 1) \Delta T$ the approximation of y_j is

$$\hat{y}_j(\alpha) = \sum_{i=1}^n b_i \phi_i(t_j, \alpha) \quad (1)$$

where α is a vector with real and imaginary eigenvalue components such that, $\phi_i(t_j, \alpha) = e^{\alpha_i t_j}$ for α_i corresponding to a real eigenvalue, and $\phi_i(t_j, \alpha) = e^{\alpha_i t_j} \cos(\alpha_{i+1} t_j)$ and $\phi_{i+1}(t_j, \alpha) = e^{\alpha_i t_j} \sin(\alpha_{i+1} t_j)$ for a complex eigenvalue.

The error between the estimate and the measurement at each time point is also known as the residual, and is found as,

$$r_j(t_j, \alpha) = y_j - \hat{y}_j(t_j, \alpha) \quad (2)$$

Another metric, known as L2 norm or Euclidean norm is used to quantify the closeness of fit of the signal reproduced from the modes i.e. $\hat{y}_j(t_j, \alpha)$ to the original sampled signal y_j ,

$$\frac{1}{2} \sum_{j=1}^m (y_j - \hat{y}_j(t_j, \alpha))^2 = \frac{1}{2} \|r(\alpha)\|_2^2 \quad (3)$$

where $r(\alpha)$ is known as the residual vector. The key unknowns to find in this problem are α and b from (1). As mentioned earlier, several methods can be used for modal analysis using signals. In this project, we have explored -

1. Variable Projection Method – It uses the Matrix Pencil method to get the initial modes.
2. Matrix Pencil Method – Research has shown that often the initial Matrix Pencil modes are sufficient [5].
3. An “Optimal” Matrix Pencil Method – A new iterative method developed in this project, to optimally select the signals to use in determining the modes to minimize the cost function (i.e. the norm in (3)). This method was extensively used in the visualizations shown further, as they significantly minimize the time required to compute and hence visualize modal information.

To find b , we use a result from the Variable Projection Method, namely

$$\hat{y}(\alpha) = \Phi(\alpha)\mathbf{b} \quad (4)$$

Then the residual is minimized by selecting,

$$\mathbf{b} = \Phi(\alpha)^+\mathbf{y} \quad (5)$$

where $\Phi(\alpha)$ is an m by n matrix (m sampled points, n eigenvalues), such that $\phi_{ji}(\alpha) = e^{\alpha_i t_j} \cos(\alpha_{i+1} t_j)$ and $\phi_{j,i+1}(\alpha) = e^{\alpha_i t_j} \sin(\alpha_{i+1} t_j)$ for a complex eigenvalue, and $\phi_{ji}(\alpha) = e^{\alpha_i t_j}$ for a real eigenvalue.

$\Phi(\alpha)^+$ is the pseudoinverse of $\Phi(\alpha)$. These steps are used with all the three modal analysis methods mentioned above, to quickly approximate each signal just by knowing α .

1.2 Signal Approximation

To reproduce a signal from its modes, the following method is used. Let n_C be the number of complex conjugate modes and n_R be the number of real modes. Let $n_C + n_R = q$. Then each signal k , $y_k(t)$ can be approximated as,

$$\hat{y}_k = f_k(t) + \sum_{j=i}^q A_j e^{\sigma_j t} \cos(\omega_j t + \phi_j) \quad (6)$$

where $f_k(t)$ is a polynomial de-trend function which is usually linear or quadratic. For real modes, $\omega_j = \phi_j = 0$. The j^{th} mode is characterized by the modal parameters: damping factor (σ_j), frequency (ω_j) and mode shape consisting of amplitude (A_j) and phase (ϕ_j). The total number of modes is q .

These methods have been tech-transferred into the latest version of PowerWorld Simulator, and can be used by BPA engineers for testing and evaluation.

The rest of Part IV shows example visualizations of modes calculated using the Optimal Matrix Pencil method. While the results are shown using simulated “PMU” data from transient stability runs, the idea is that these techniques can be applied to actual PMU measurements as well.

2 Data Inputs and Analysis

2.1 Procedure for Determining the Modes in Software

First, we briefly touch upon the method to calculate these modes and their associated quantities such as angles, magnitudes, etc. in order to visualize them. The signals used in this modal analysis can either be from 1) actual grid measurements such as PMU data in a variety of formats such as COMTRADE, csv, JSIS files, or 2) simulation results obtained from a transient stability run. For a meaningful analysis the input data should have at least some oscillations; these typically occur in response to a disturbance event. Figure 1 shows the selection of simulated PMU data from transient stability results.

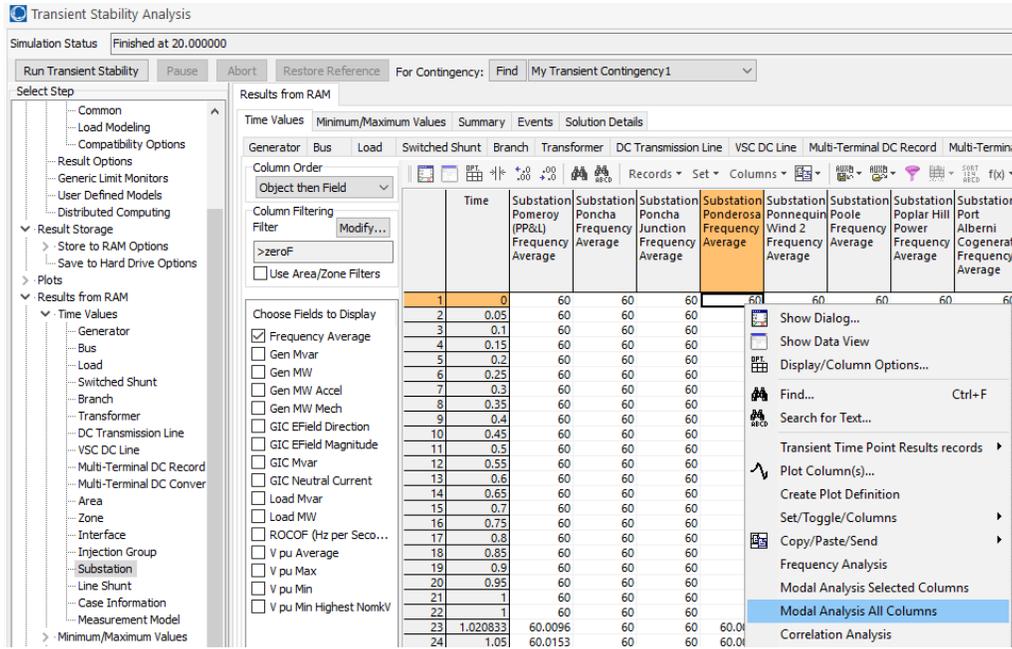


Figure 1. Transient stability results window for modal analysis

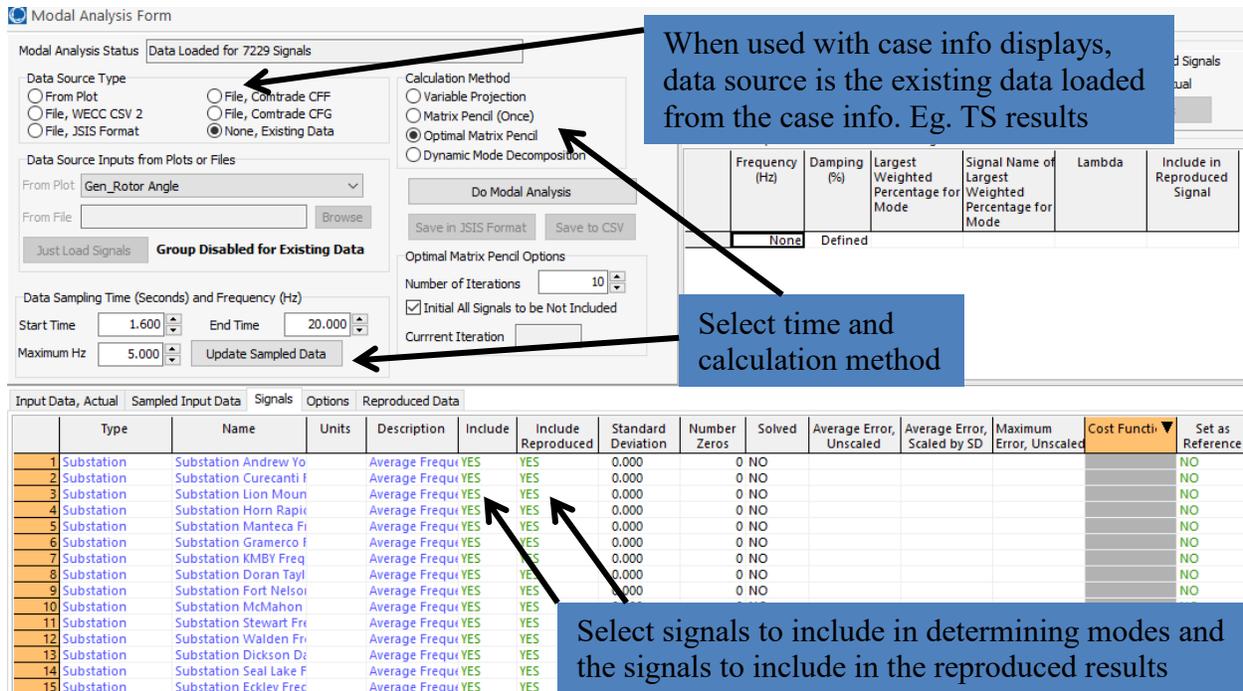


Figure 2. Modal analysis form for selecting inputs, sampling, method, and other parameters

Figure 2 briefly describes the modal analysis window. If the Optimal Matrix Pencil Method is chosen (as was done in this report), one does not need to explicitly specify the signals to include. They will be included optimally based on the highest cost function, at each iteration. The number of iterations can be changed by the user; by default this field is set to 10. In our experience, choosing between 10 and 15 iterations has produced reasonable results. The first modal analysis run can be done with the default 10 iterations, and this parameter can be increased later based on the highest cost function and the closeness of the corresponding reproduced signal to the original.

Figure 3 shows the results after performing modal analysis with 10 iterations. These results correspond to a solid three-phase bus fault contingency applied at t=1 sec., cleared by opening the lines connected to the bus at t=1.6 sec. The analysis is then performed on the average frequency calculated at each substation, depicted in Figure 13. We rank the cost function in descending order, and then right-click and open the dialog for the substation associated with the highest cost function. In this case, it is Substation Seven Mile. The dialog box shows values for the original signal, the signal reproduced from the estimated modes, and their difference. Plotting the first two yields Figure 4. It shows that while the reproduced signal captures the overall variation of the original, it under-represents the amplitude of the high frequency oscillation. Since the reproduced signal is generated from the estimated modes, this means that our estimates are off by that measure. Looking at Figure 3, we see that the original signal was not include in the estimation by the optimal matrix pencil algorithm. In order to improve the estimate, we can include this signal by toggling this to YES, increasing the iteration number by 1, and running the modal analysis again. Note that these calculations are really fast, taking about 25 seconds for ~7200 signals with 191 time points and 10 iterations (using i7-7820HQ CPU @ 2.9 GHz, 32 GB RAM).

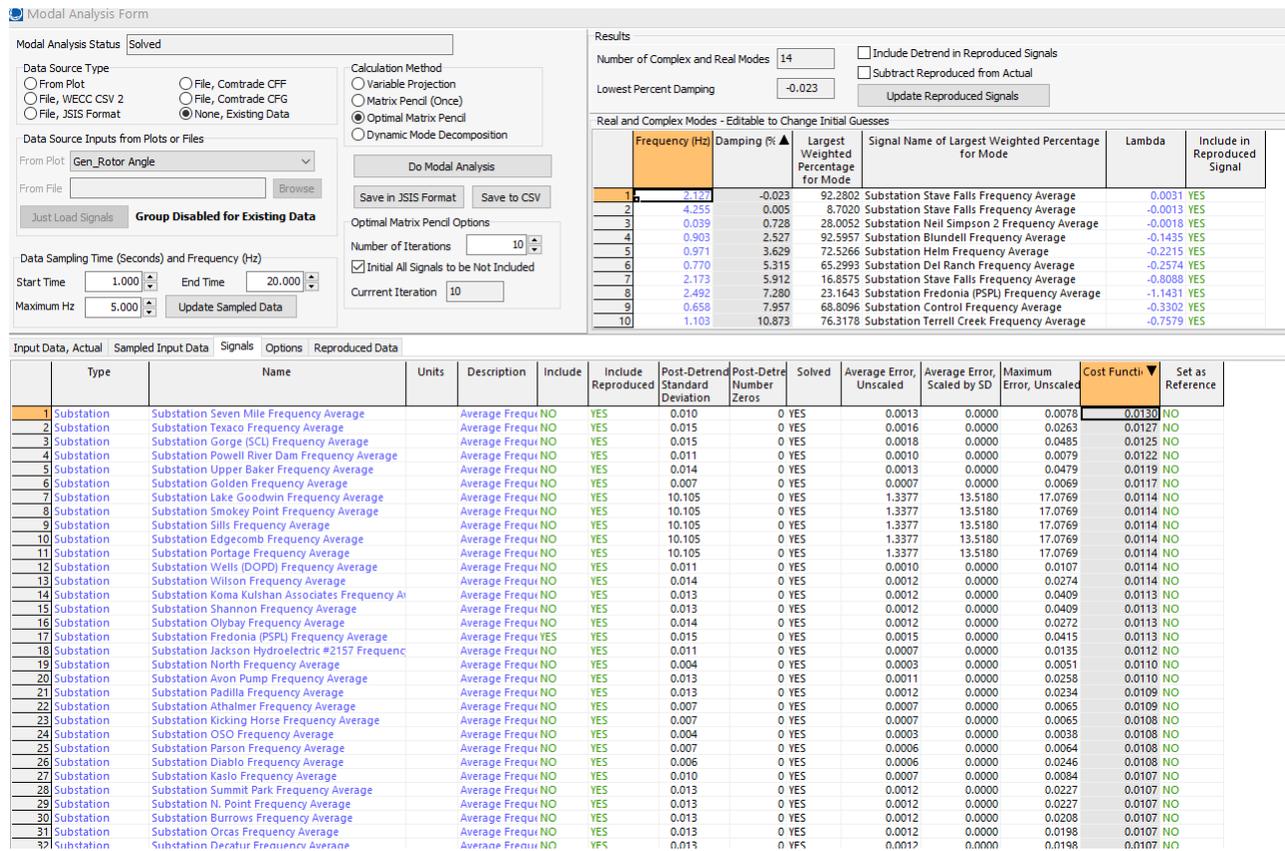


Figure 3. Example modal analysis results for a WECC transient contingency. Signals used are the average of the bus frequencies at each substation.

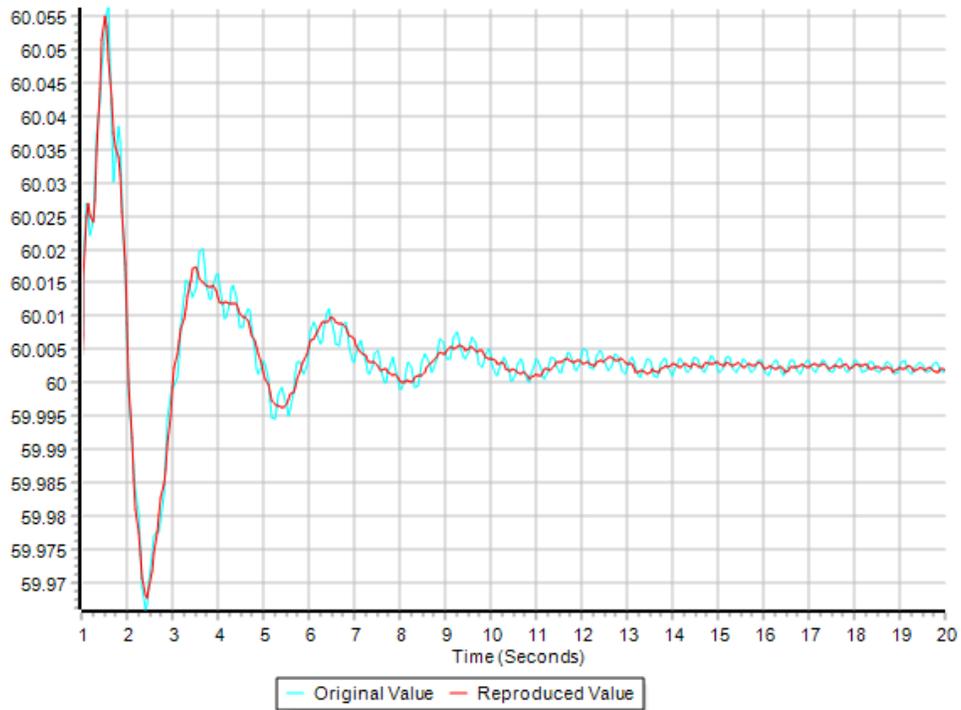


Figure 4. Comparison between the simulated (original) average frequency at Seven Mile substation, and the frequency reproduced by the calculated modes. This substation corresponds to the highest cost function at the end of the first modal analysis run, with 10 iterations i.e. only 10 original signals included.

Figure 5 is the result of including the signal and re-running the modal analysis, increasing the number of iterations by 1. The cost function for the Seven Mile substation frequency average reduces by a factor of almost 2.

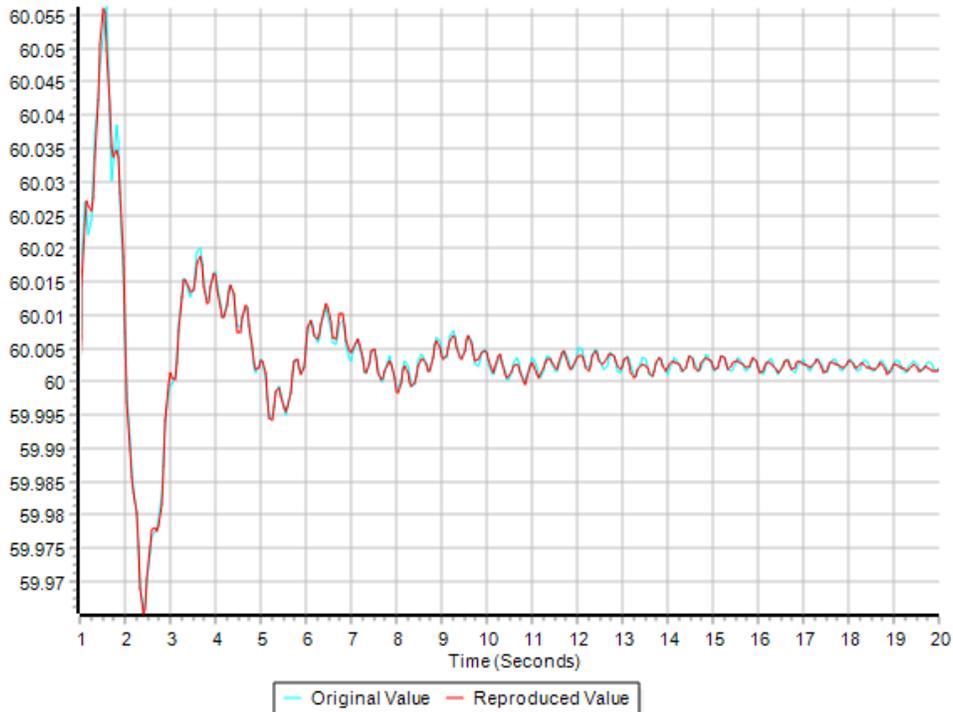


Figure 5. Updated results at Seven Mile substation when its original signal is included in the analysis

Looking at the overall modal analysis results, the largest cost function has now decreased to 0.0126 from 0.013. Figure 6 shows the comparison between the original and reproduced signals for the new signal with the highest cost function (Substation Gorge SCL frequency average). The original signal is much smoother compared to the reproduced, and the former was again not included in the analysis. To improve the estimate, Figure 7 shows the Substation Gorge SCL average frequency comparison with slight improvements, after including the original signal in the analysis. Note that at this stage the number of iterations has increased to 12. The cost function for this signal is now 0.0088, with the highest function value now reduced to 0.0110.

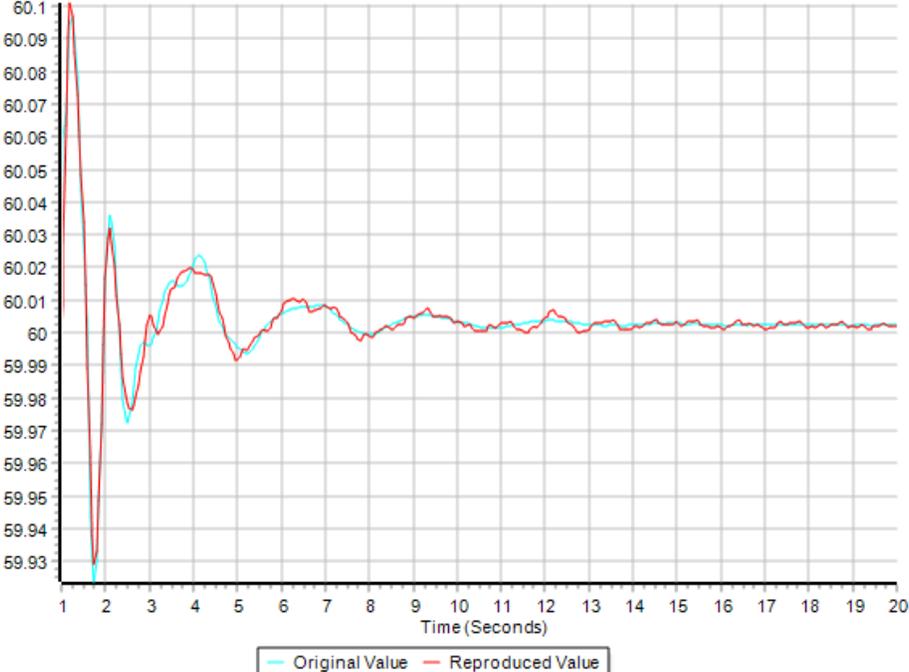


Figure 6. Average frequency at Substation Gorge SCL. This is the signal with the highest cost function value after 11 iterations.

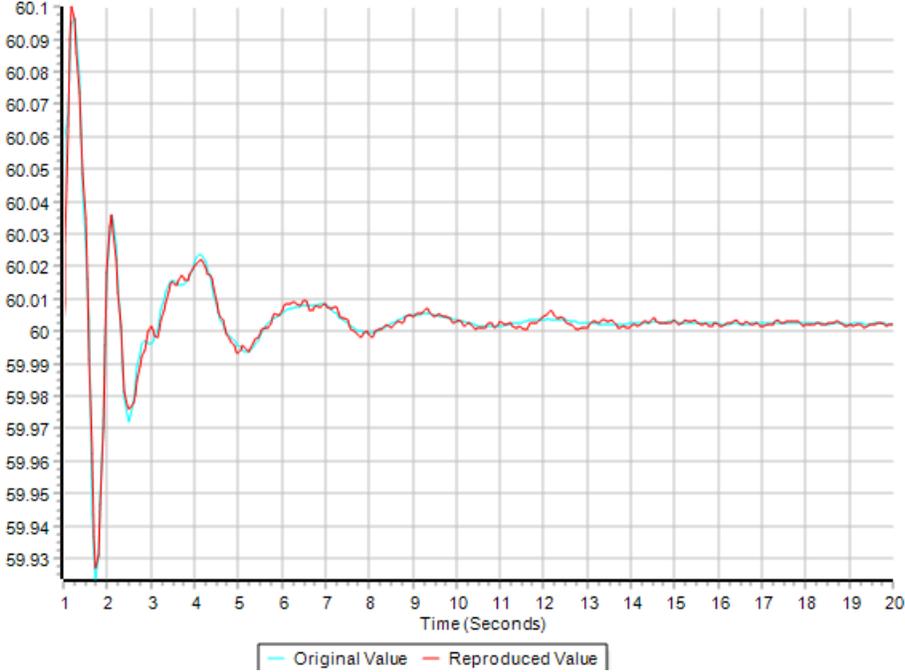


Figure 7. Results after including the original Substation Gorge SCL average frequency signal

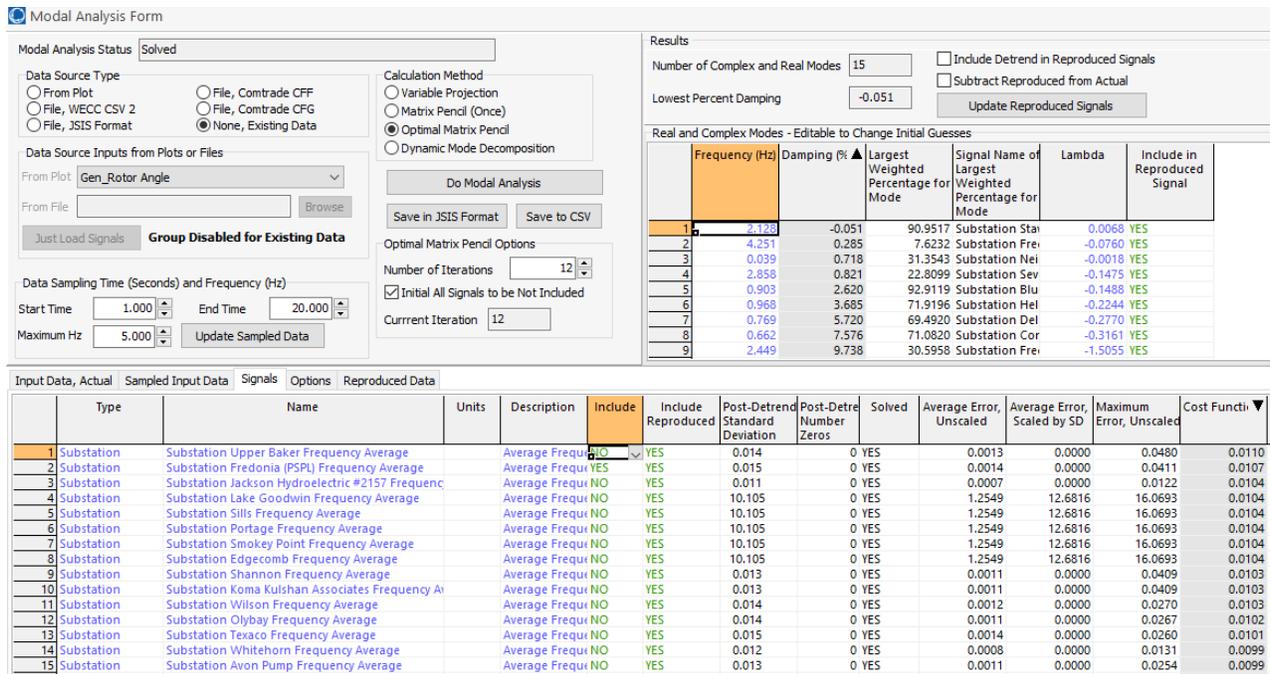


Figure 8. Final modal analysis results for substation average frequencies

Figure 9 shows the cost functions at each substation for this modal analysis run. If this is acceptable, including the worst case scenario of Figure 10 (cost = 0.011), we can move on to further analyzing the individual modes.

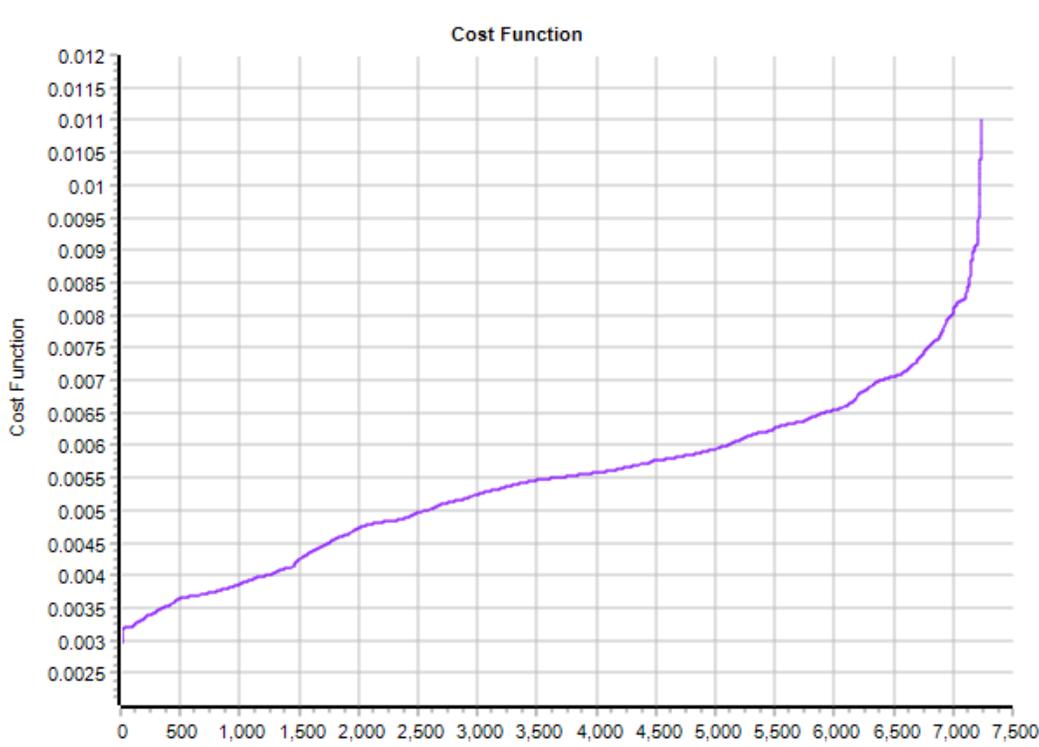


Figure 9. Cost function of all signals ranked in ascending order

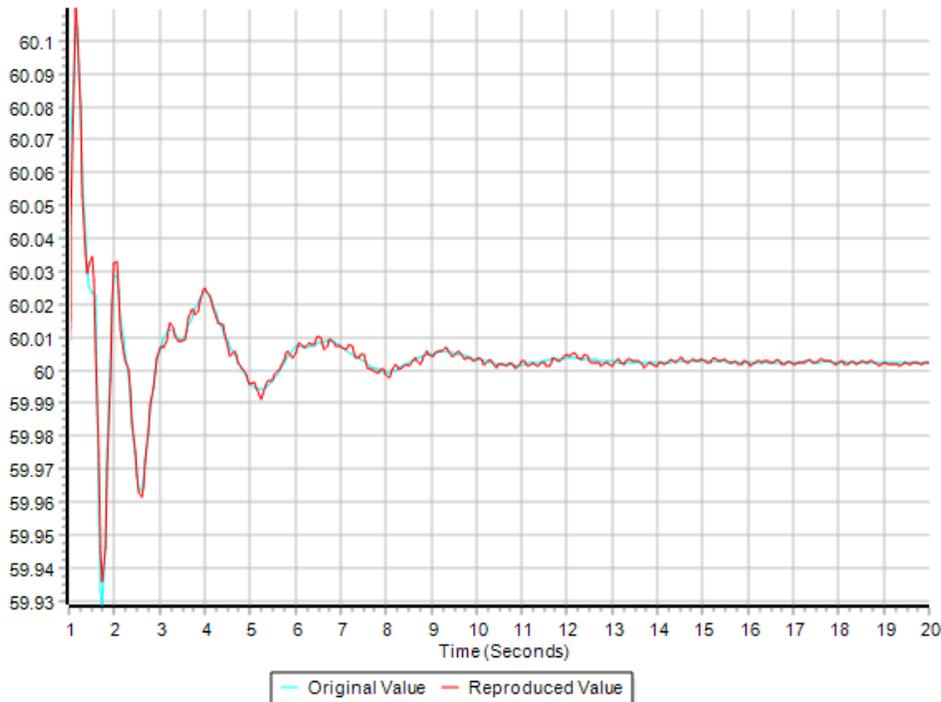


Figure 10. Upper Baker substation average frequency comparison i.e. maximum cost function signal at the end of 12 iterations

2.2 Effect of Windowing

In a real-time environment with continuously streaming PMU data, the analysis will need to be done “online” i.e. near real-time if not real-time. The incoming data will need to be analyzed in pieces, or in other words, windows with certain number of time points. Also, it is probably impractical to have the modal analysis engine running continuously during ambient conditions, and should be triggered when an event is detected. Here we discuss the effects of these parameters on the modal analysis results.

1. Start Time of the Window – Figure 3 shows that the start time for the modal analysis was chosen to be $t = 1$ sec. This is time point when the fault was applied. Note that between $t = 0$ to 1 sec., no disturbance was applied, with the fault cleared at $t = 1.6$ sec. Now if we repeat the estimation of modes starting from $t = 2$ seconds, we can expect to get slightly different results (with the same end time as the previous case). We also keep the number of iterations the same, i.e. 12, between the two scenarios for consistency.

The highest cost function now becomes 0.0152, implying slightly worse results compared to the $t = 1$ to 20 seconds scenario. This can be explained by the fact that there are several switching events that occur between $t = 1$ and 1.6 seconds, which are a part of the contingency. These contribute to the dynamics observed up to 2 seconds. Including them improves the estimates to an extent. Figure 11 also shows changes in the damping values of some of the modes, and slight changes in their frequencies.

However, note that when the no disturbance (i.e. flat) period between $t = 0$ and 1 second was included, the maximum cost function actually increased compared to when the start time was $t = 1$ second. This means that the start time of the modal analysis window should be as close to the inception of the disturbance (here $t = 1^+$), and avoid including too much ambient data before the disturbance. This makes accurate and prompt event detection from streaming PMU data important for achieving good modal analysis results.

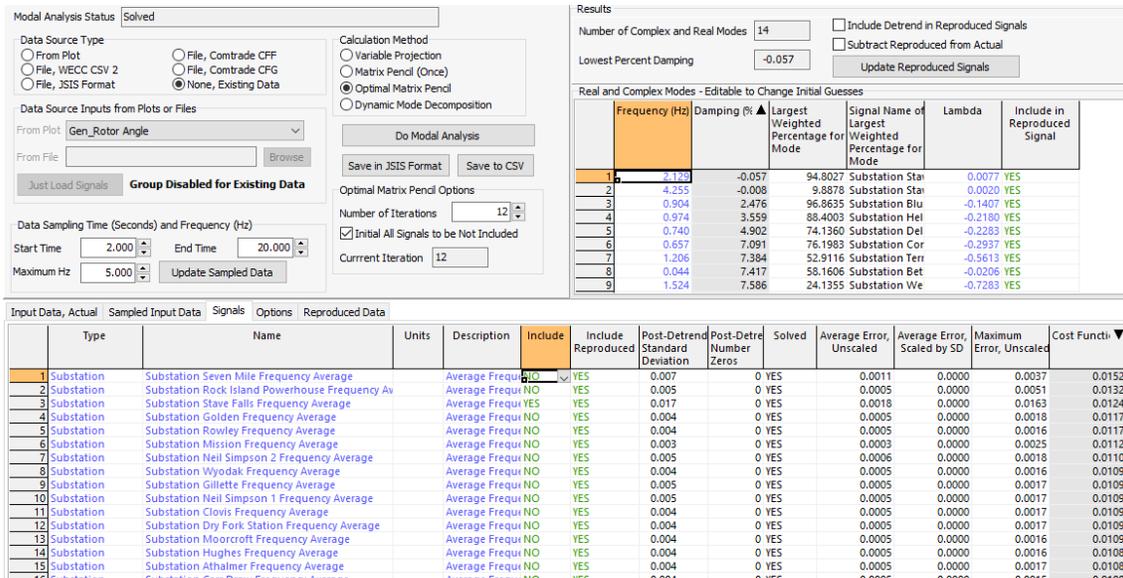


Figure 11. Modal analysis results with a different start time but same end time and number of iterations as before. Note the increase in the cost function values.

2. Duration of the Window – Comparing Figure 12 where the window is shortened to end with $t = 15$ seconds, to Figure 8, the highest cost function value increases by around 25%. We can also notice differences in the damping values of some of the modes. For example, the 4.26 Hz mode goes from being slightly positively damped to very slightly negatively damped. Negative damping implies growing oscillations, whereas a damping of 0% represents sustained, undamped oscillations. The higher the value of the positive damping percentage, the quicker the oscillations die out. In the example here, the discrepancy arises since we neglect the dynamics of this mode towards the end of the simulation. This result is, of course, dependent on the start time as discussed earlier. In real-time applications, one cannot select a very large window for analysis, since it will not only increase computation time but one has to also wait for that amount of time to elapse to process the PMU data collected during it. Hence a balance has to be achieved between the timeliness of obtaining results in a real-time environments, with their increased accuracy associated with larger windows.

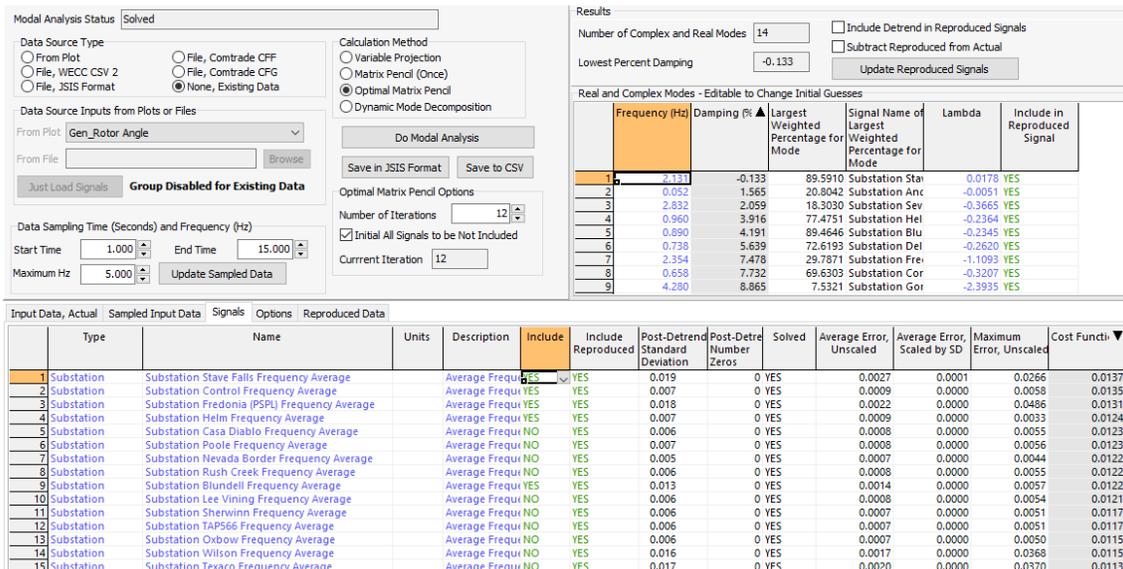


Figure 12. Modal analysis results for a smaller time window (end time is $t = 15$ seconds instead of 20)

3 Visualization of Modes - Example Scenarios

3.1 Frequencies

Figure 13 shows the average of the bus frequencies at approximately 7200 “live” substations in the 15hw WECC planning model. A 3-phase bus fault was applied at $t = 1$ sec. and cleared after 0.6 seconds, leading to sustained oscillations. The fault was cleared by opening lines connected to the bus. From a glance, the system appears to have some well-damped lower frequency oscillations, and higher frequency undamped oscillations.

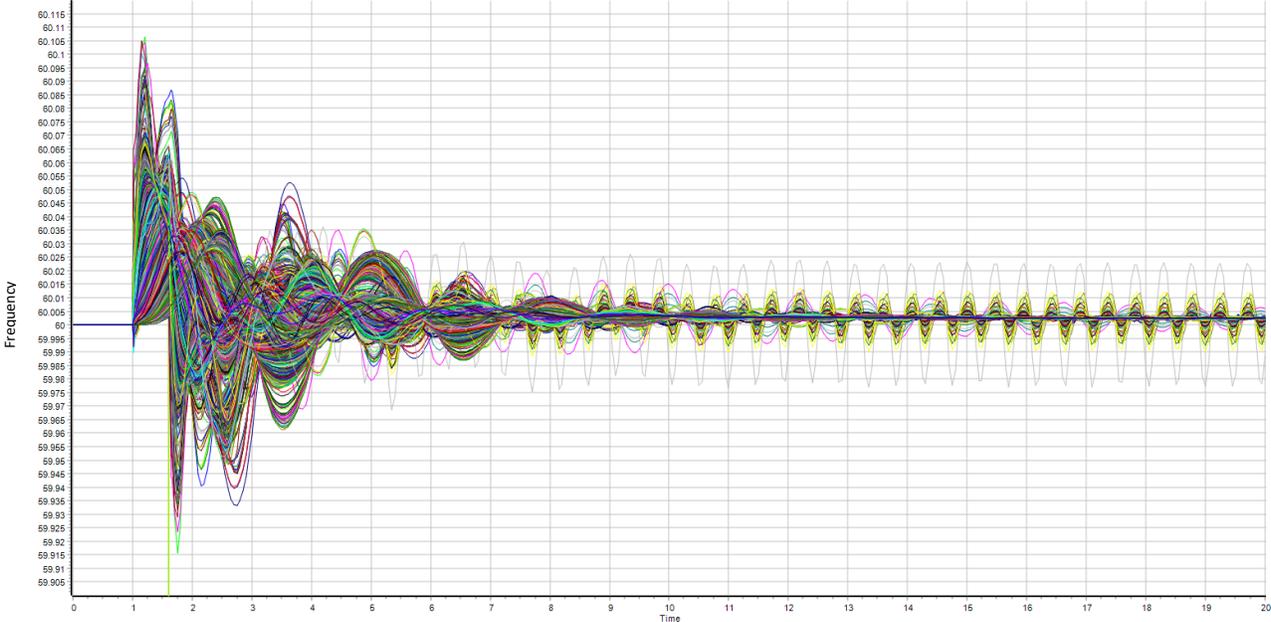


Figure 13. Substation average frequencies in the WECC system for a transient contingency

Table 1. Modal analysis results in ascending order of damping. The known WECC modes are highlighted.

	Frequency (Hz)	Damping (%)	Largest Weighted Percentage for Mode	Signal Name of Largest Weighted Percentage for Mode	Lambda
1	2.128	-0.051	90.9517	Substation Stave Falls Frequency Average	0.0068
2	4.251	0.285	7.6232	Substation Fredonia (PSPL) Frequency Average	-0.0760
3	0.039	0.718	31.3543	Substation Neil Simpson 2 Frequency Average	-0.0018
4	2.858	0.821	22.8099	Substation Seven Mile Frequency Average	-0.1475
5	0.903	2.620	92.9119	Substation Blundell Frequency Average	-0.1488
6	0.968	3.685	71.9196	Substation Helm Frequency Average	-0.2244
7	0.769	5.720	69.4920	Substation Del Ranch Frequency Average	-0.2770
8	0.662	7.576	71.0820	Substation Control Frequency Average	-0.3161
9	2.449	9.738	30.5958	Substation Fredonia (PSPL) Frequency Average	-1.5055
10	1.119	9.773	73.5678	Substation Terrell Creek Frequency Average	-0.6906
11	2.229	11.692	20.1606	Substation Stave Falls Frequency Average	-1.6487
12	0.360	14.412	83.8328	Substation Andrew York Frequency Average	-0.3298
13	0.726	14.544	58.8148	Substation Gorge (SCL) Frequency Average	-0.6704
14	1.404	16.027	38.1938	Substation Gorge (SCL) Frequency Average	-1.4328
15	0.250	31.891	56.2616	Substation Madden Frequency Average	-0.5288

It is evident from Figure 13 that there are several modes of oscillation, and it is difficult to distinguish between them visually. This is why modal analysis is necessary to identify the individual components, as shown in Table 1. For the rest of this section we use signals sampled between $t = 1$ and 20 seconds, 12 iterations, and simulation results from Figure 13, unless otherwise specified.

There are well-known modes of oscillation in the WECC system [6], which have also been highlighted in Table 1. These are -

- “North–South Mode A” nominally near 0.23 Hz. This was historically termed the “NS Mode.” Its properties are well known. In this, Alberta swings against the rest of the system, while British Columbia (BC) and Pacific Northwest (PNW) swing with Alberta.
- “North-South Mode B” nominally near 0.4 Hz. This was historically termed the “Alberta Mode.” Its properties are well known. Here, Alberta swings against BC and Northern WECC which swing against Southern WECC. This is the most wide-spread mode in the system.
- “East-West Mode A” nominally near 0.45 Hz. Until 2013, this mode was not observed due to poor PMU coverage.
- “British Columbia” mode nominally near 0.6 Hz. The properties of this mode are fairly well known.
- “Montana” mode nominally near 0.8 Hz. Its properties are well known.

Another helpful point to know is that, in BPA’s oscillation detection tools and methods, high energy oscillations are typically studied in four frequency bands [7]:

1. 0.01 to 0.15 Hz band –for governor, plant controller or AGC control oscillations
2. 0.15 to 1 Hz band – for electromechanical inter-area and some local plant oscillations, and also oscillations caused by sustained operation of a hydro-power plant in rough zone
3. 1 to 5 Hz band –for local generator oscillations, generator excitation controls, PDCI controls, etc
4. 5 to 15 Hz high frequency band – for steam-generator torsional, PDCI, wind generator control oscillations, sub-synchronous interactions.

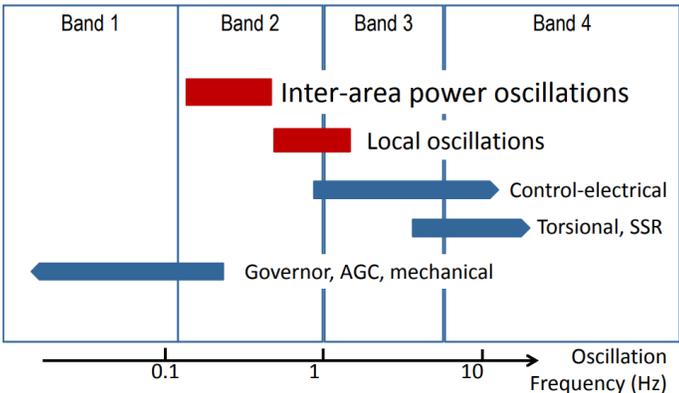


Figure 14. Frequency bands for analyzing oscillations in WECC

Figure 14, reproduced from [7] depicts these bands; there is a clear overlap in some of them. The WECC modes and oscillation frequency bands were discussed here to better explain, and even validate the modal analysis results and visualizations. For the visualizations, an information layering technique through the use of geographic data view (GDV) objects and contouring is used to present large amounts of modal and other processed results [8][9][10]. Information pertaining to the wide-area oscillation mode activities and source of oscillations are used to characterize the system dynamics.

The main quantities we visualize are, a) Mode angle, b) Mode magnitude, and c) Cost function. This is repeated for different measurements such as substation average voltage magnitudes, and bus voltage angles.

For illustration, we begin with one of the known modes, i.e. the 0.25 Hz mode, also known as the NS A mode. This mode is well damped in this case, i.e. 31.89% damping to be precise (see Table 1). The detailed results associated with this mode can be viewed, as shown in Figure 15. Figure 16 shows all the reproduced signals, accounting for this mode only. This can also be referred to as the mode components. It is seen that the mode is infact well damped , with the oscillations halting at around 10 seconds. Another way to interpret this figure would be, is that this is the deviation caused in the substation average frequencies due to this mode of oscillation.

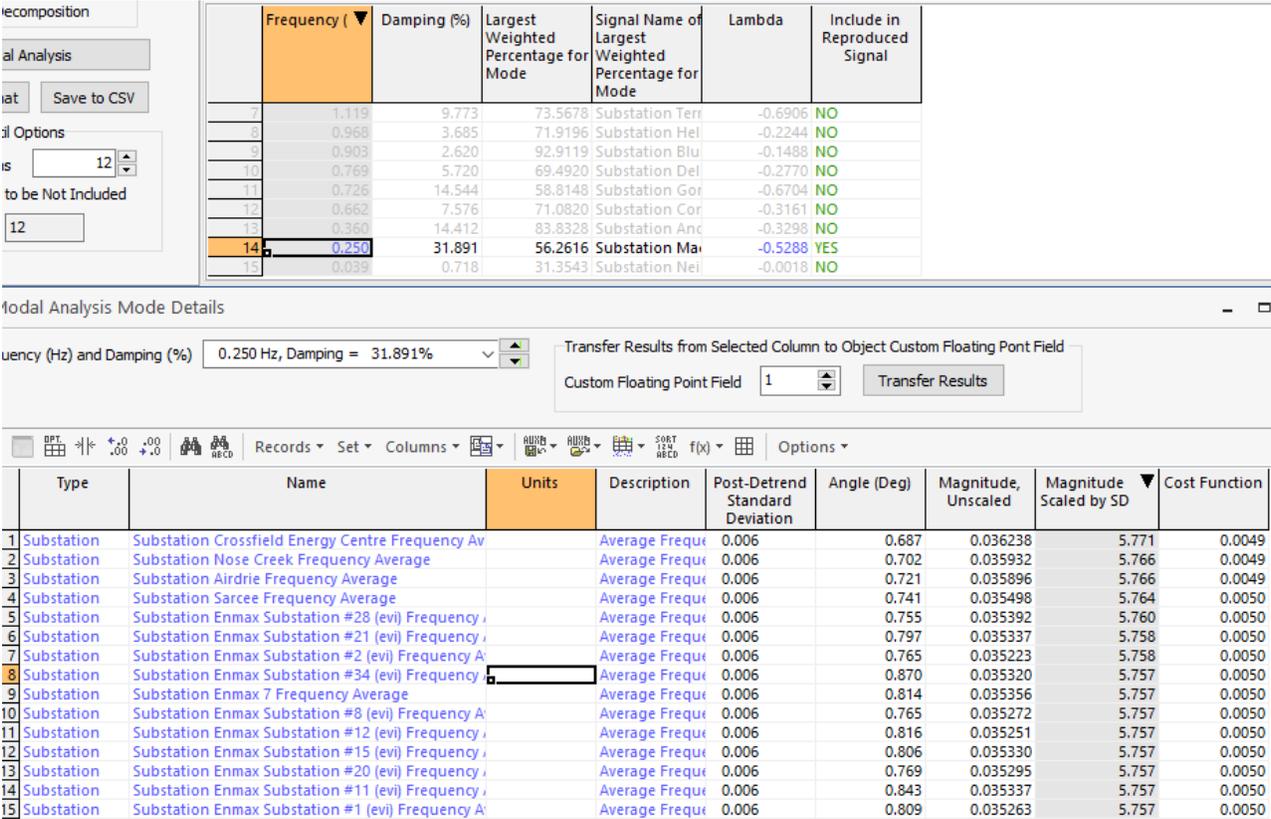


Figure 15. Detailed results for the 0.25 Hz mode

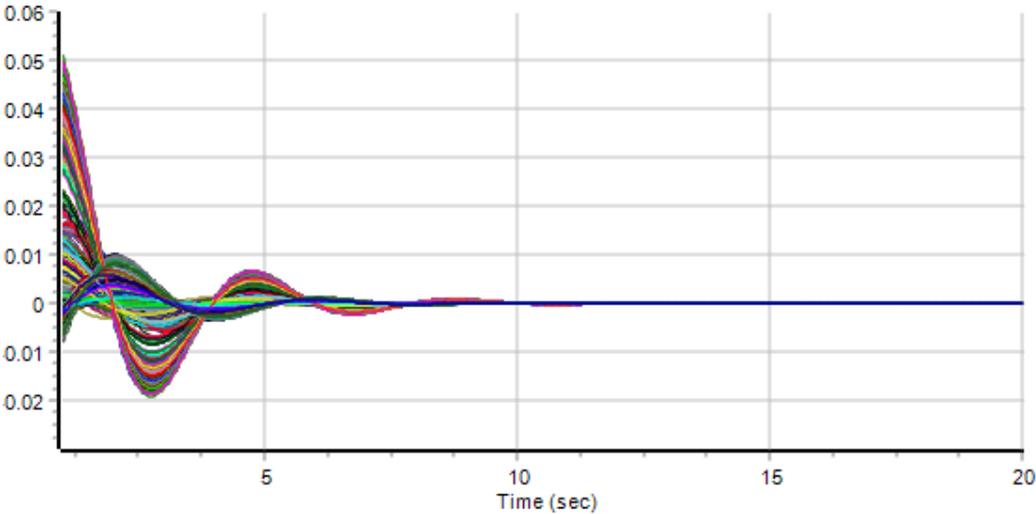


Figure 16. Frequency deviation (mode component) at each signal (substation average frequency) due to the 0.25 Hz mode

We first visualize the mode angles. These are a key element of mode shape visualizations. The angles are denoted in degrees and vary between -180 and +180 degrees, with the idea that the ends of this range indicate the same value, following a circular convention. To contour this, a color map should be chosen such that as the angles move closer towards ± 180 degrees, they should converge to the same or similar color. A new, cyclic color map was implemented for this very purpose. This discrete, 12 color map has each division representing 30 degrees. Discrete maps can also be modified into continuous maps by specifying this in contour options.

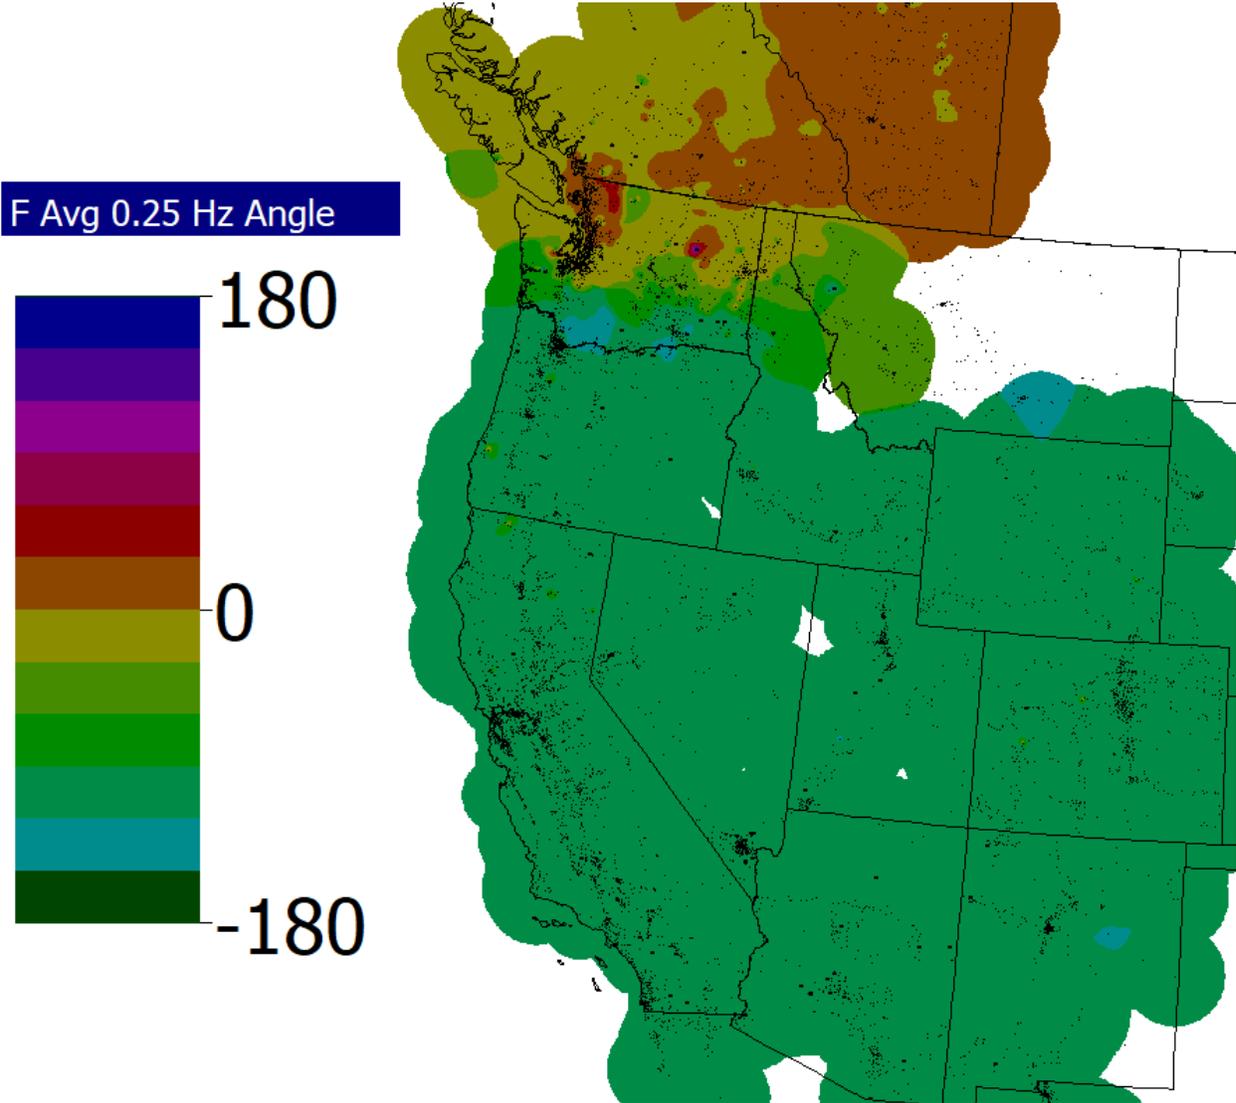


Figure 17. 0.25 Hz mode angles

The angles contoured in Figure 17 indicate that the Alberta region, with an angle of around 0 degrees, oscillates against the rest of the WECC contoured in teal green (around -120 degrees). British Columbia and some parts of Washington also oscillate with Alberta (-30 to 30 degrees). This agrees with the description of the NS A mode mentioned earlier, thus validating the modal analysis methodology and calculations.

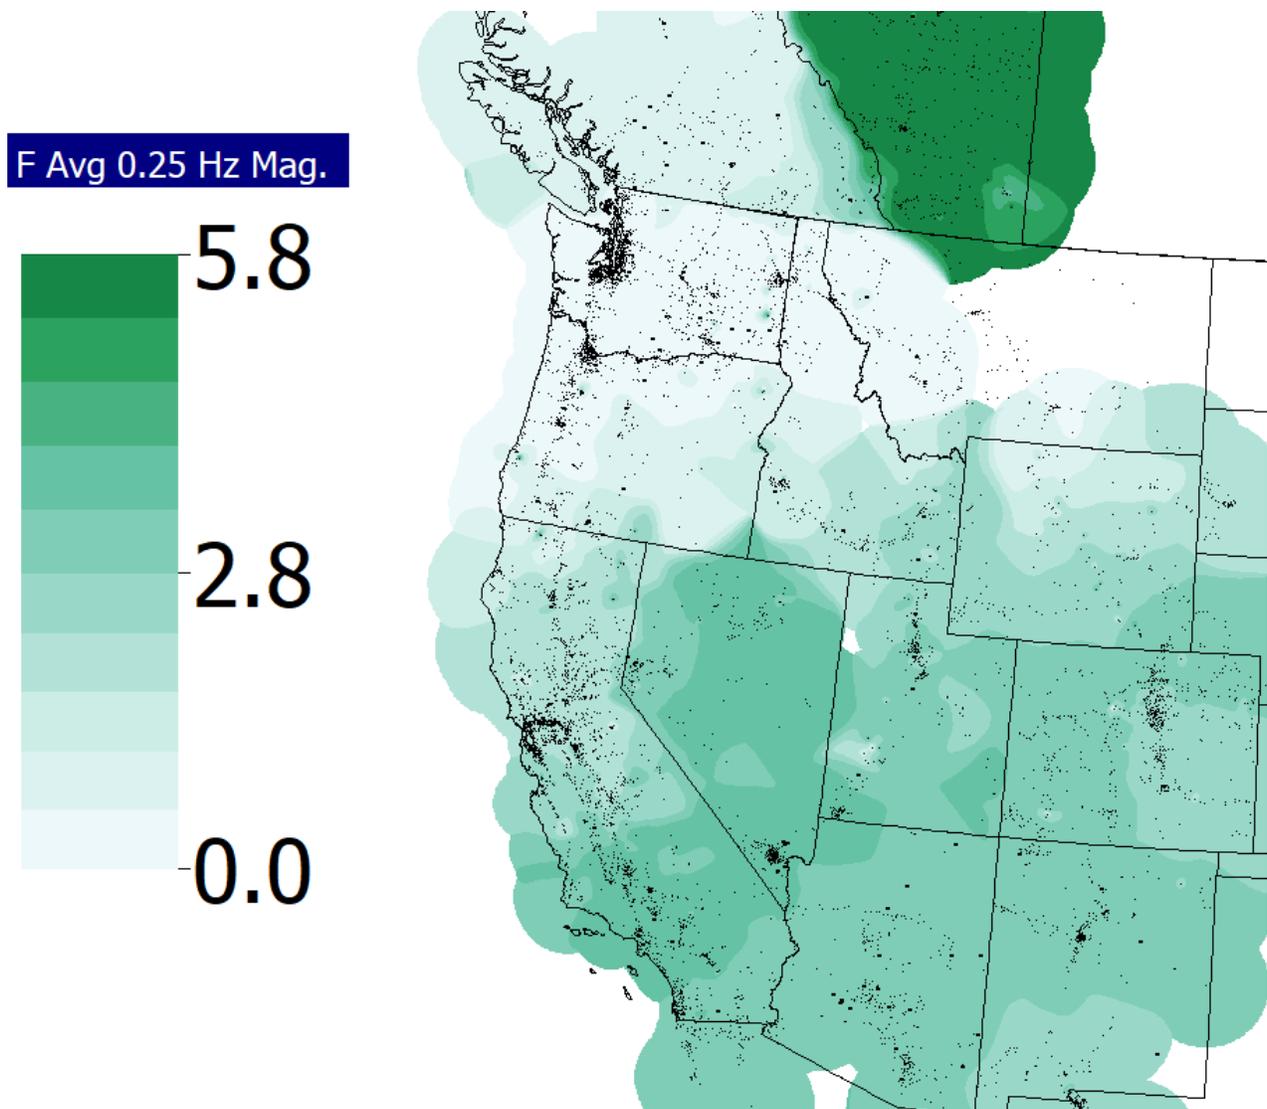


Figure 18. 0.25 Hz mode magnitudes (scaled by standard deviation)

Figure 18 shows the magnitudes associated with this 0.25 Hz mode, for each of the substation frequency average signals. All the magnitude contours in this report have been scaled by standard deviation. One can hypothesize the oscillation / transfer of energy between a small area such as Alberta and the remaining part of WECC (discounting pacific North West, swinging almost with Alberta), which are both ends of the network, to be causing the magnitude distributions for this mode as seen here.

Similarly, we visualize the remaining known WECC modes, i.e. 0.36 Hz, 0.66 Hz and 0.77 Hz, found in our simulations results.

For the 0.36 Hz mode, the oscillations seem almost completely out of phase with another, as seen in Figure 19. This should help delineate the two areas of the “inter-area” oscillation fairly easily. Also we see from Table 1 that the damping for this mode is less than that of the 0.25 Hz mode.

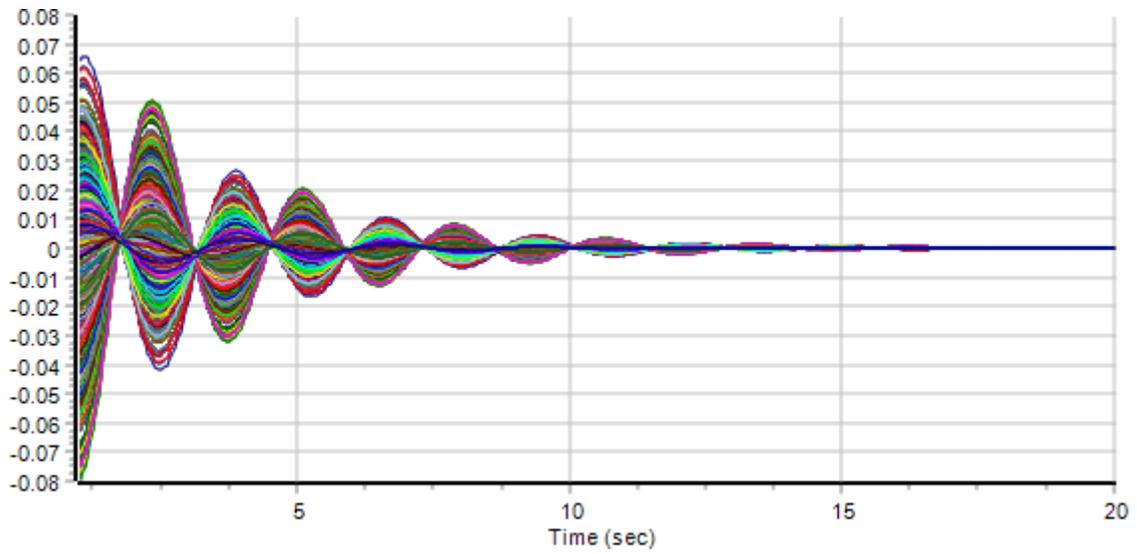


Figure 19. 0.36 Hz mode components or frequency deviations for each signal

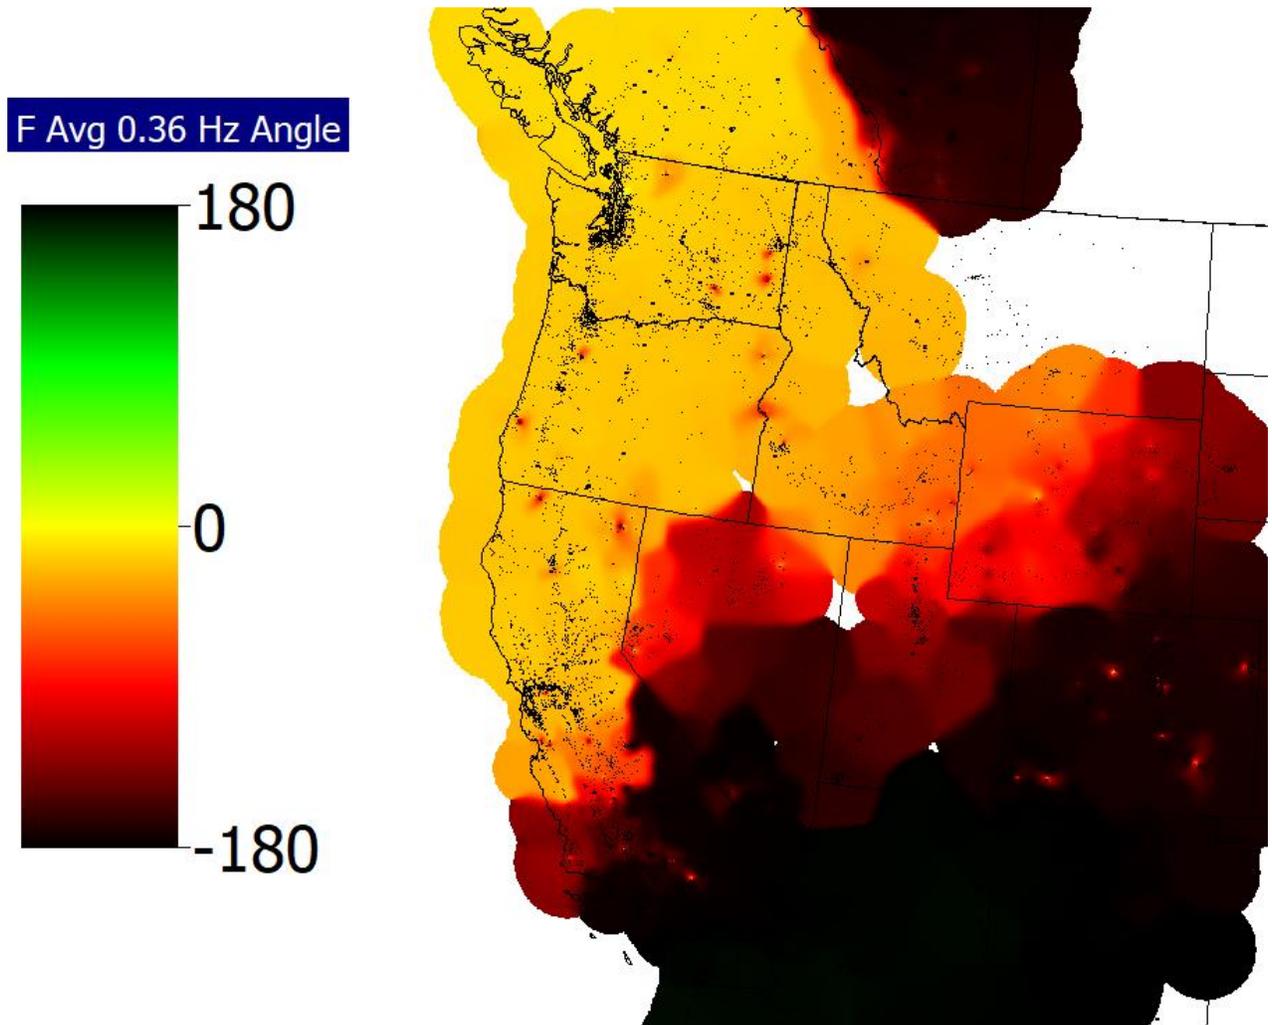


Figure 20. 0.36 Hz mode shape using a cyclic color map (Cyclic RGB and Black)

We use another color map in Figure 20 to show the mode angles. This can be used when the signal components are almost anti-phase (as in Figure 19), hence several colors are not needed to interpret angle values. The angle distribution in Figure 20 matches the description of the NS B mode.

Figure 21 shows the contour of the magnitudes.

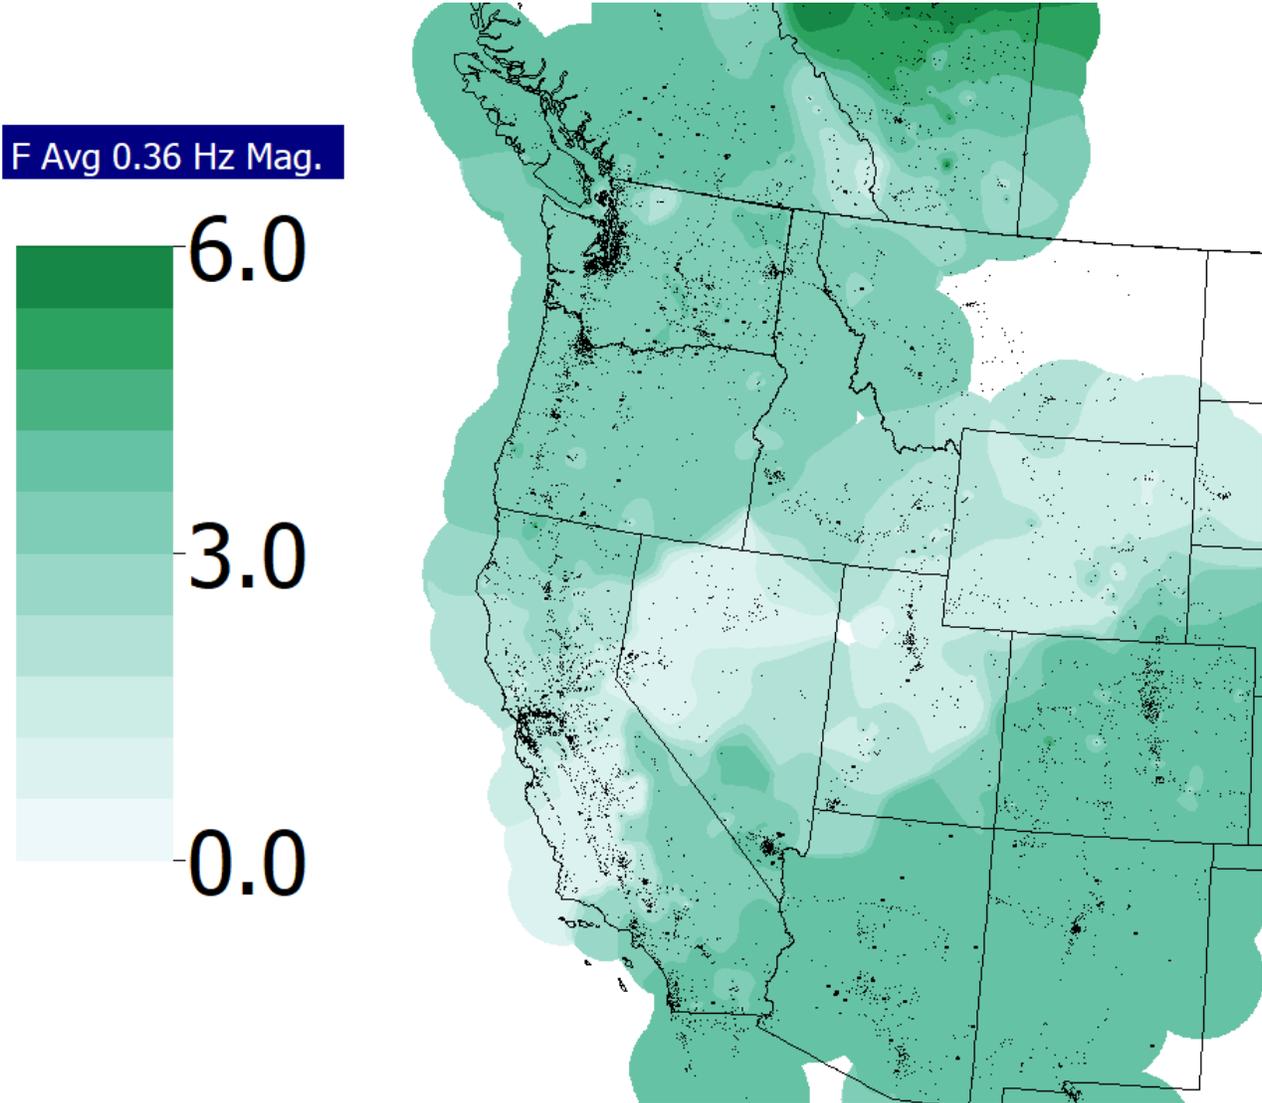


Figure 21. Magnitudes for the 0.36 Hz mode

For strong inter-area oscillations such as the 0.25 Hz and 0.36 Hz modes, the angle contour can provide sufficient information about which areas are oscillating in unison with and against which areas. The magnitude contours can help determine the regions where frequency (in this case) deviates the most.

Next we look at some of the other WECC modes. Figure 22 and Figure 23 shows details of the 0.66 Hz mode. This falls in the vicinity of the British Columbia (BC) mode. One example of the BC mode visualization is shown in Figure 40 of [11]; it is a 0.68 Hz oscillation for a 2012 heavy winter model as opposed to the 2015 heavy winter model used in our case. Comparing that to Figure 23 below, it is clear that in both of them, BC oscillates against WA and surrounding areas strongly, i.e. the phase difference is almost 180 degrees. BC also is almost in phase with NV. When we say one area oscillates with or against the other, we actually refer to the

generators located in those areas oscillating in or out of phase with each other.

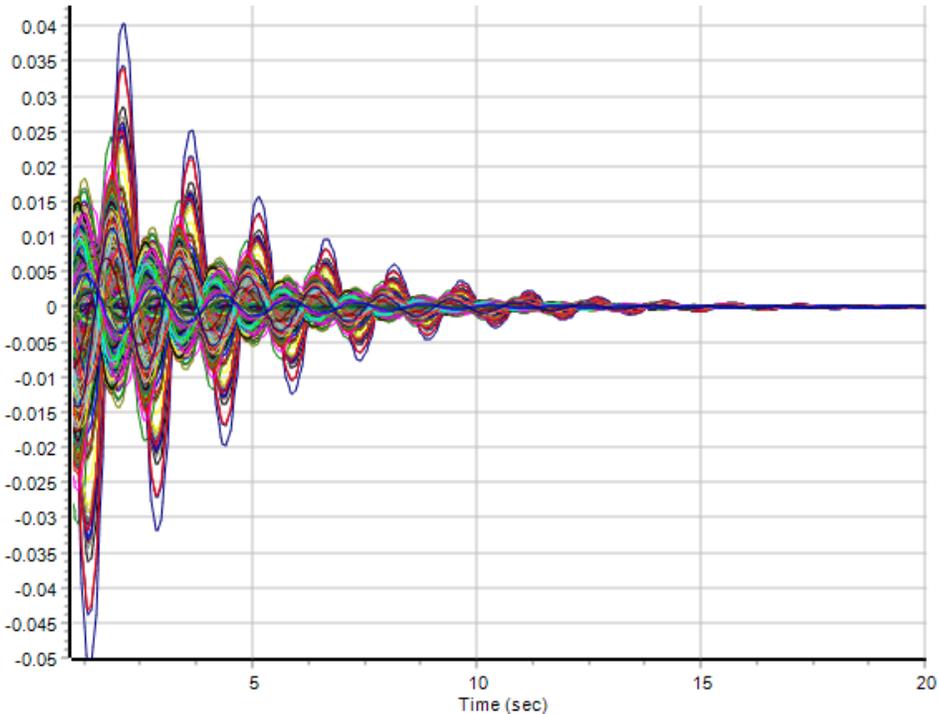


Figure 22. 0.66 Hz mode components

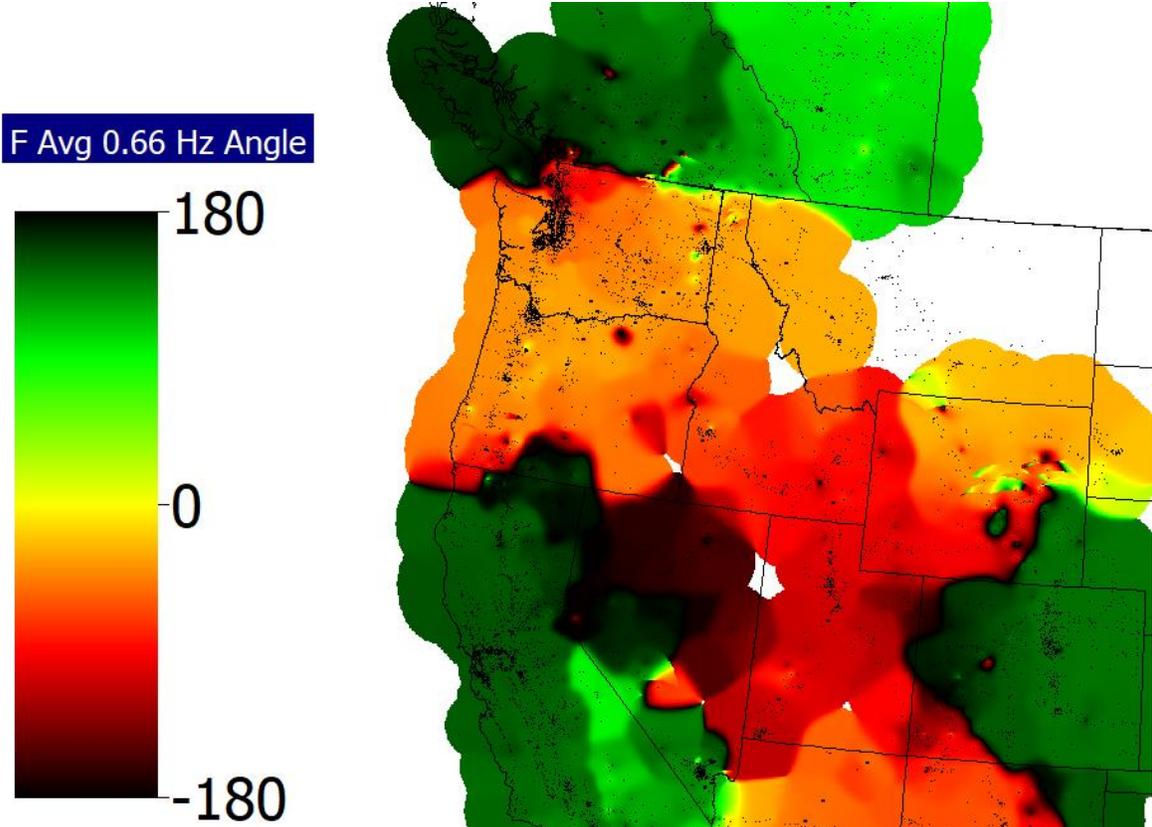


Figure 23. 0.66 Hz mode angles

Similarly, we visualize the 0.77 Hz mode, referred to as the Montana mode. Again referring to Figure 43 of [11] we see that MT and WY oscillate together against the lower part of WECC, including CA, NV, and UT.

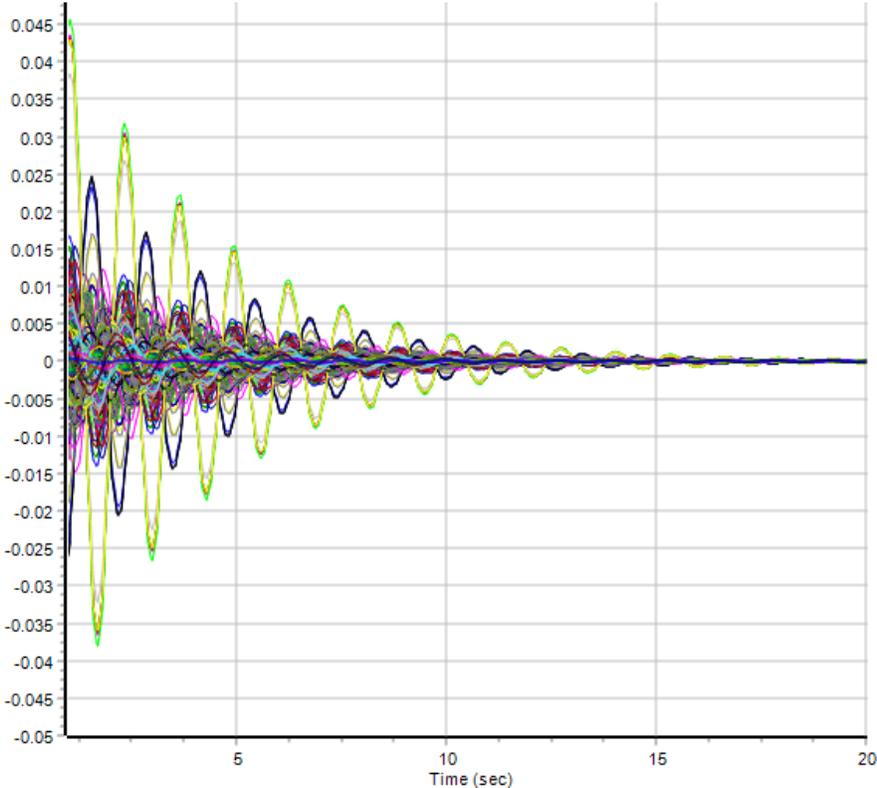


Figure 24. 0.77 Hz mode components

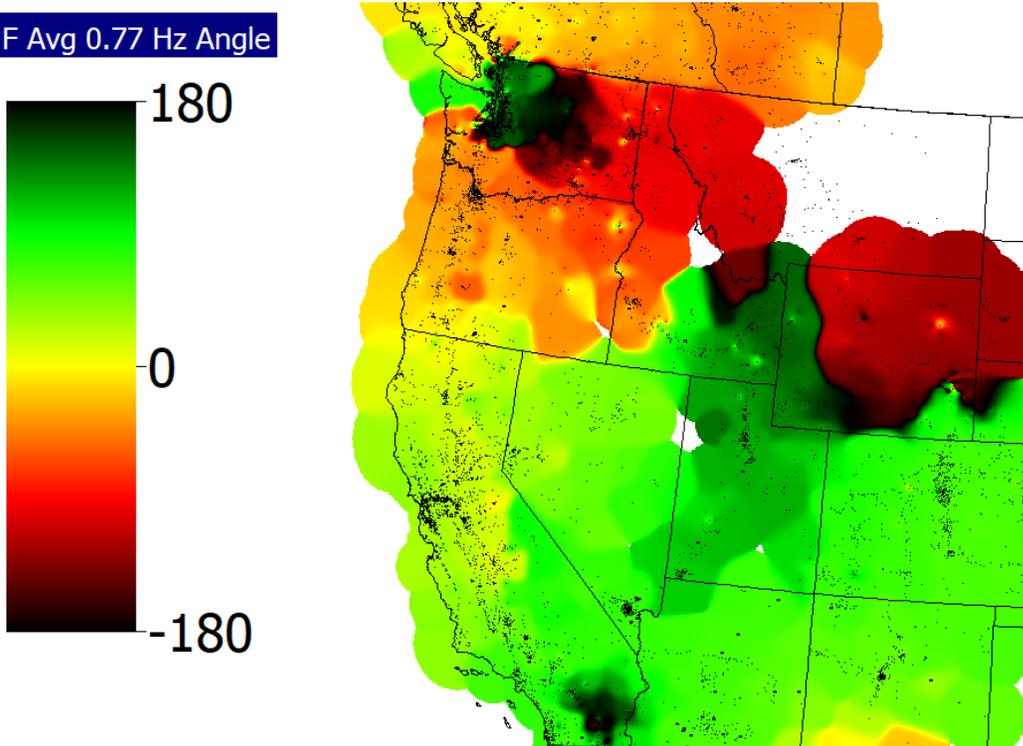


Figure 25. 0.77 Hz mode angles

In the previous two examples, while the angle contours were helpful in distinguishing between the areas of the inter-area oscillations, there was information lacking on the magnitudes. In case of the Montana mode, Figure 24 in particular shows that there are a few generators where the oscillation is very pronounced. The magnitude information would be more useful and stand out in this case, compared to the 0.25 Hz and 0.36 Hz oscillations. Later in the report, we will discuss how we can visualize magnitudes and angles on the same display effectively, instead of having to look at two different images.

Next we contour the cost function associated with these results. Note that the cost function is independent of the individual modes, and depends on the modal analysis solution.

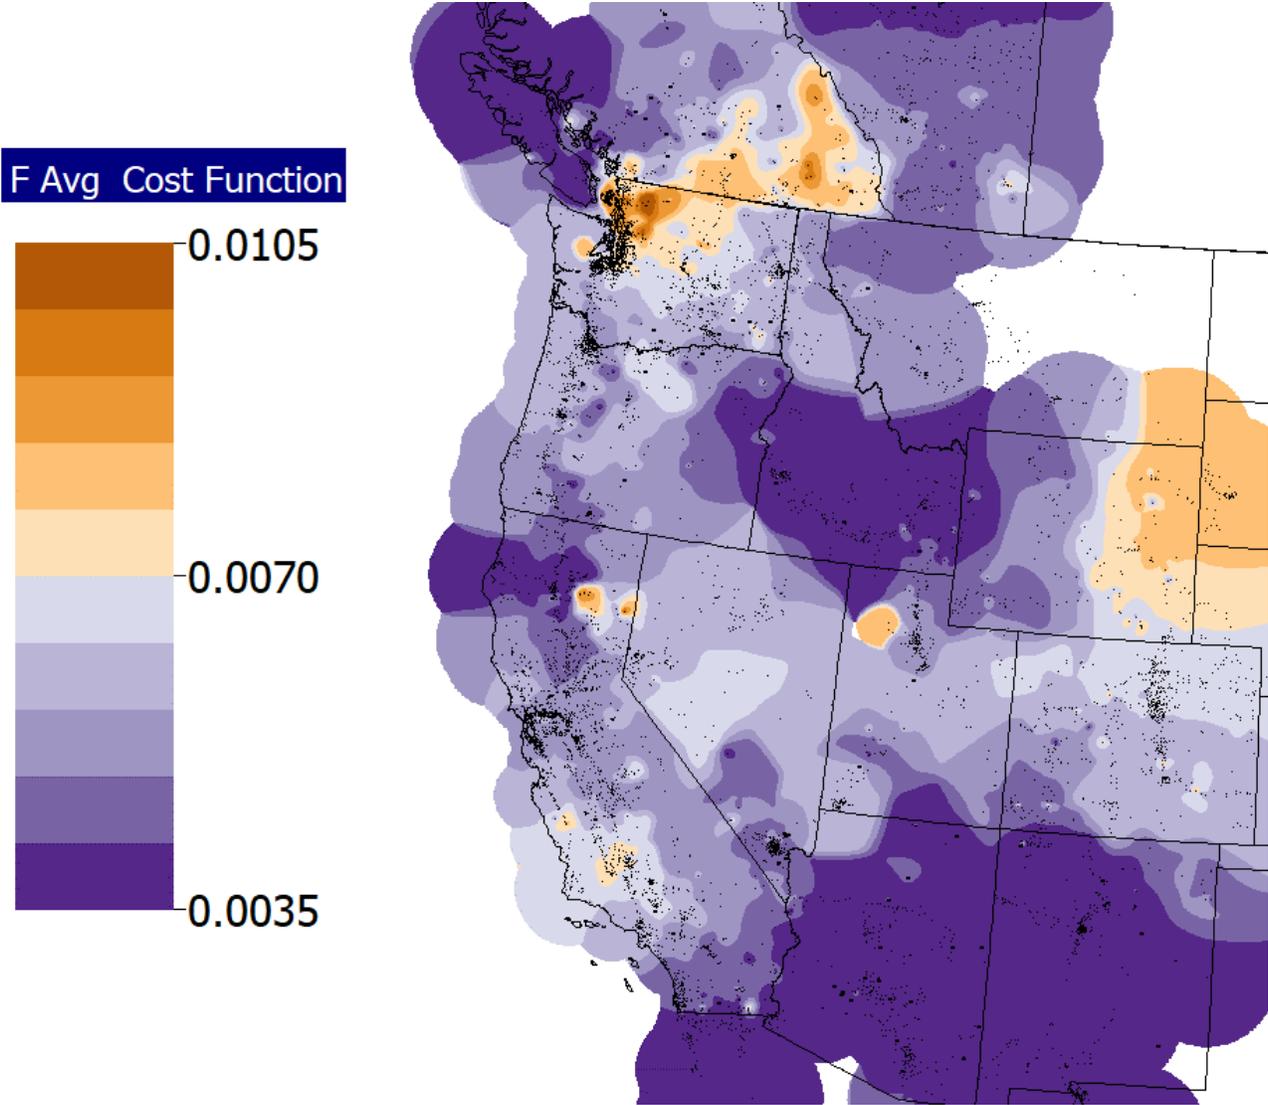


Figure 26. Cost function for all signals in the modal analysis run

The lower the cost function, the better is the estimate of the signals in that part of the system. This color scheme was chosen because the colors can quickly separate the “good” zones from the “not so accurate”. This distinction would depend on the threshold value selected (here 0.007) that separates the two colors. The user may choose this threshold through heuristics, or after some offline studies. The software also suggests these ranges based on calculations such as the maximum, minimum and median values. Figure 26 shows the modal estimates are most accurate in ID, AZ and NM, as well as some parts of BC, Alberta, and OR, among the northern WECC states. This does not necessarily mean that the estimates in the yellow/brown zones are incorrect. The threshold we selected here is very low. This is just an example of how the cost function can be

visualized.

Going back to Table 1, we now consider the mode with the least damping. This is a 2.13 Hz mode which is in fact slightly negatively damped (-0.05%). Figure 27 shows one of the signals which has a major component of this mode.

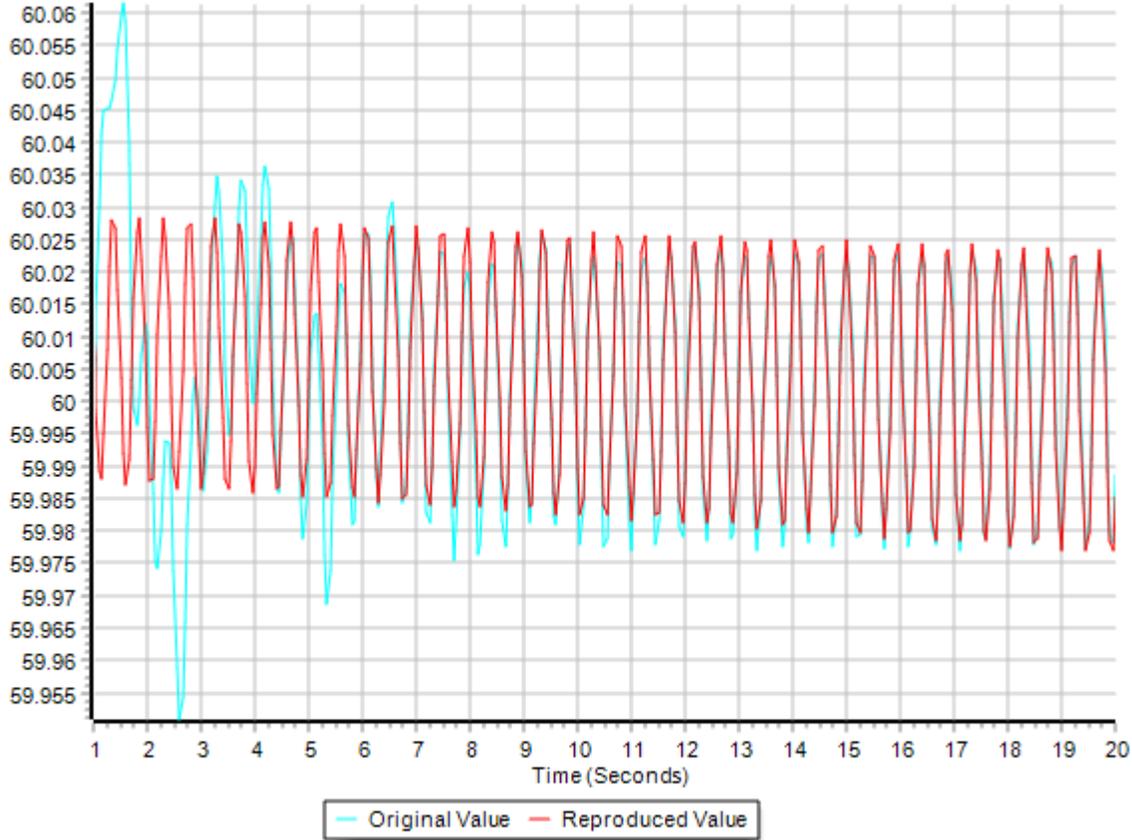


Figure 27. Substation Stave Falls average frequency

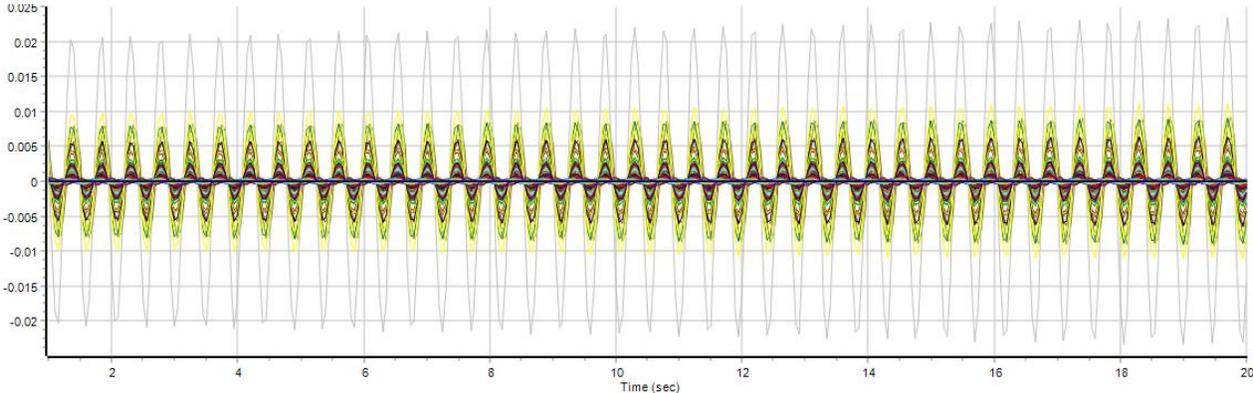


Figure 28. 2.13 Hz mode components

Next we visualize the angles of this mode.

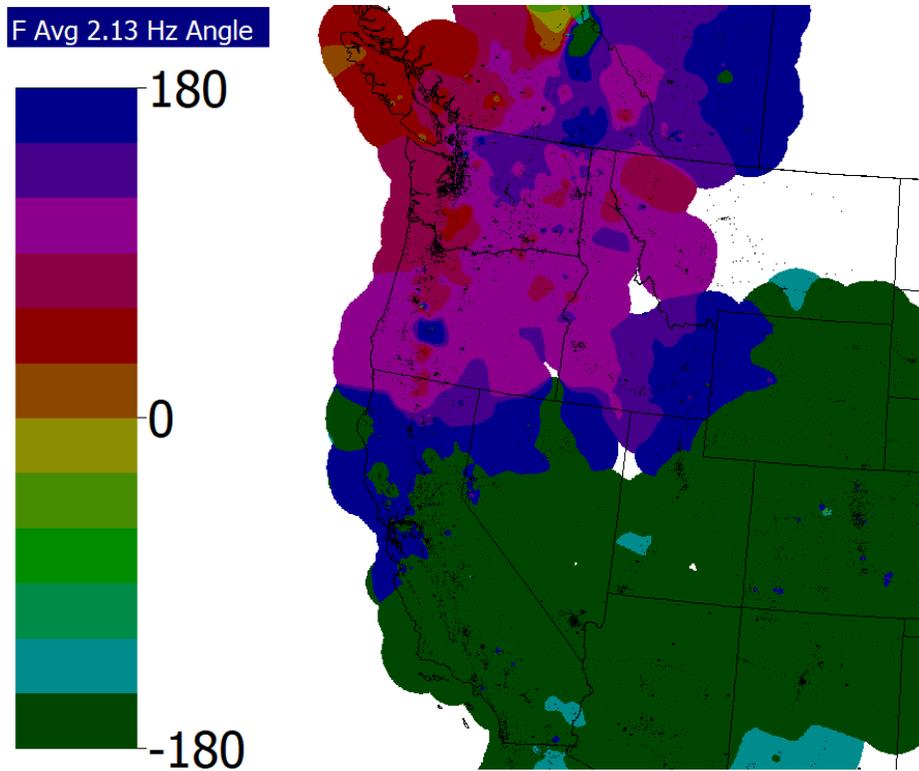


Figure 29. 2.13 Hz mode angles

In some situations such as this one, it may not be overtly obvious to tell whether it is an inter-area or local plant oscillation just by looking at the snapshot of the angles (although the frequency bands discussion indicates this is most likely a local generator mode). We look at the magnitudes as well to gain more insight.

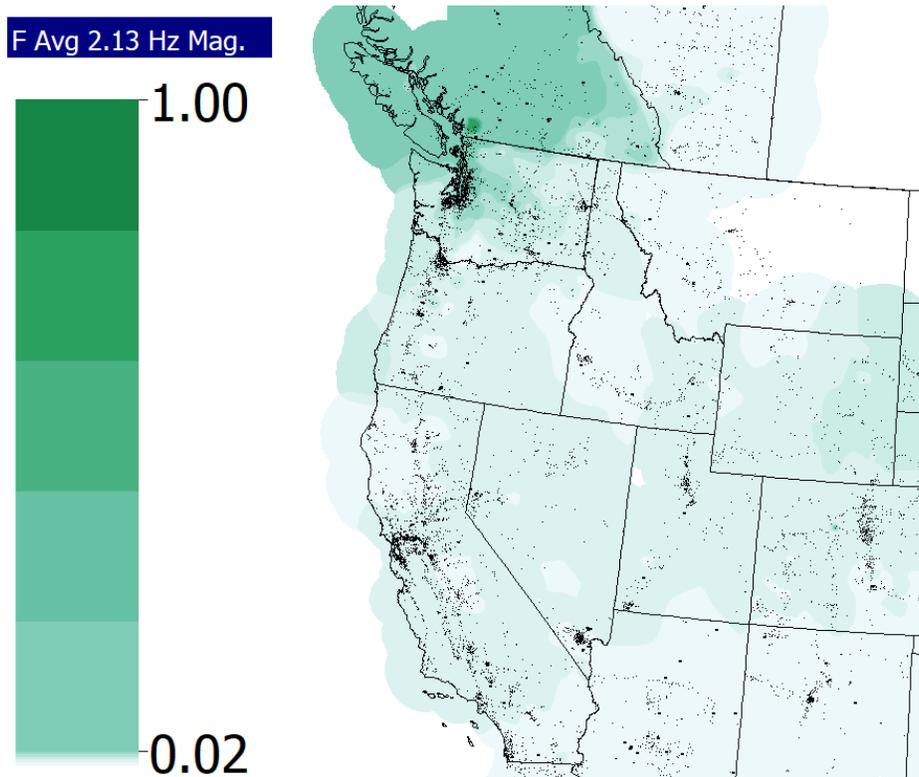


Figure 30. 2.13 Hz mode magnitudes

Figure 30 shows the magnitude contour for this mode. The location with the largest magnitude is in British Columbia, indicated by the small, dark green dot just above the US-Canada border, namely the Stave Falls substation. A hypothesis is that this oscillation may be originating at or close to this substation.

As briefly mentioned earlier, magnitudes and angles can be shown on the same graph to convey a more complete picture of the system, without obscuring critical information or overwhelming the user. One such example is of the angles being denoted by arrows, as in Figure 31. These arrows are defined under “Geographic Data Views”. The value of the angle is represented by the direction of the arrow. The arrows could have uniform size and thickness, or vary with the mode magnitude. Arrows allow the user to visualize the angle variations more discretely. The resolution of arrows can be increased as per user settings beyond what is shown here; Figure 31 is a “pruned” version of the GDV of angles at all substations (which does look crowded), for better visualization.

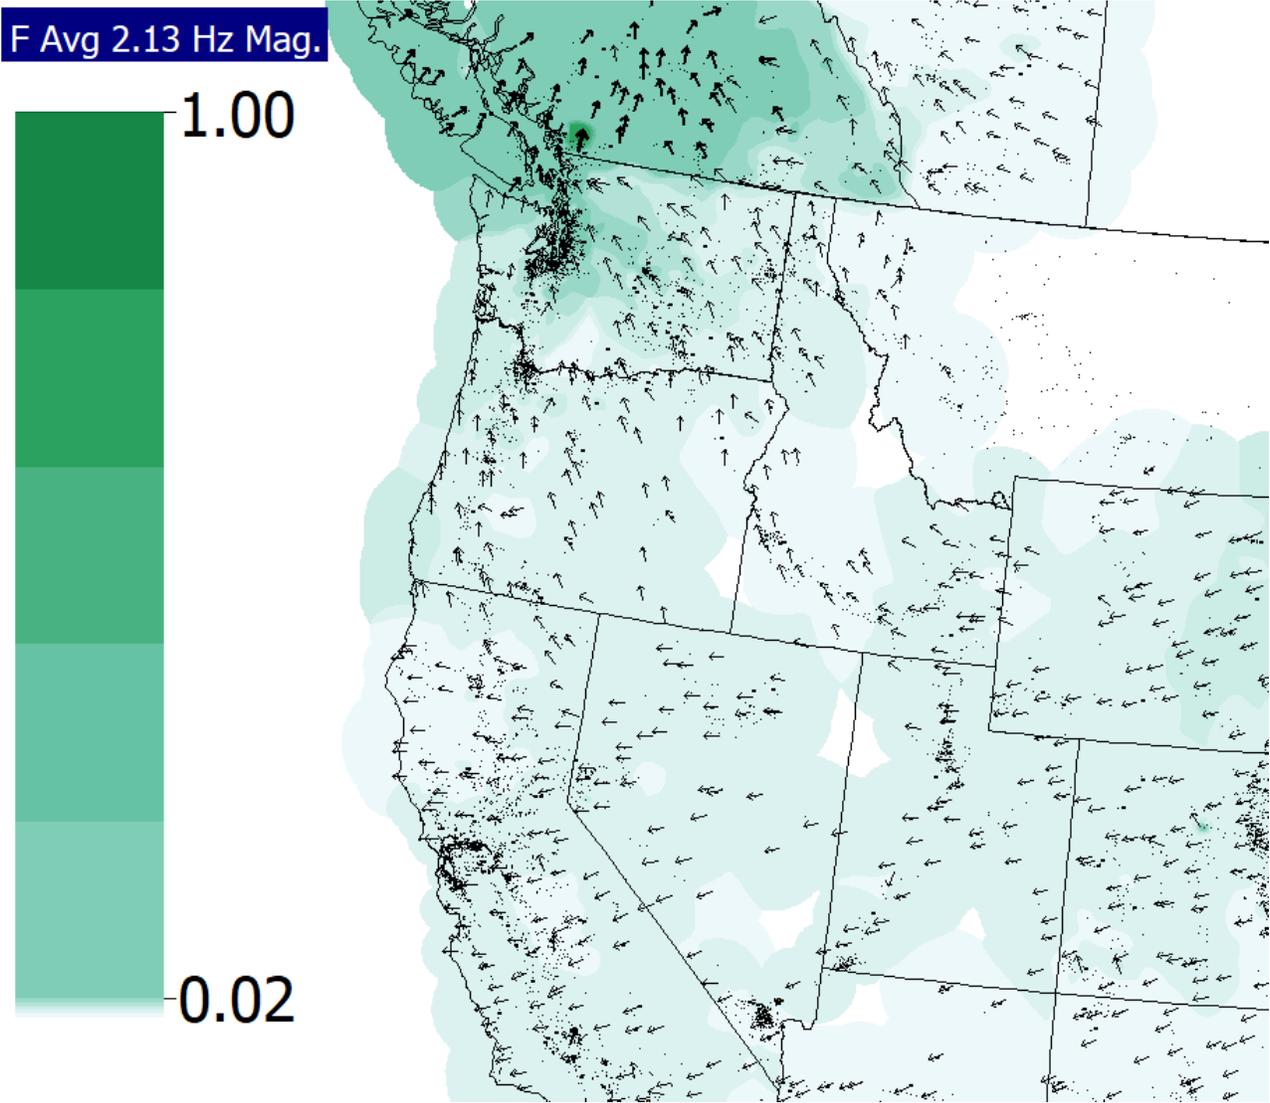


Figure 31. 2.13 Hz mode magnitude contoured, and angles shown by arrow directions.
Size (thickness) of arrows is proportional to the mode magnitudes

The dark thick arrow, indicating a large magnitude, coincides with the small, dark green dot region in BC captured by the contour, thus showing the effectiveness of visualizing angles this way. While this was a straightforward example, sometimes a GDV can help better distinguish between substations to find potential sources of oscillations. Note that GDV and contours are independent visualizations, and one is not necessary

for, nor dependent on the other. A user can choose to visualize mode magnitude using just the contour or just the GDV. The same applies for angles. Arrows can be used to represent both angle and magnitude.

Both the magnitude contour as well the arrow GDV hint to the fact that Substation Stave Falls is a possible source of the oscillation. This will be confirmed in the next section in which we implement a method to detect the source of oscillations. It is also interesting to note here that in Table 1, for 2.13 Hz mode, 90.9% of it is contributed by a single signal which happens to be the Stave Falls substation average frequency.

If local angle variations are of interest, the arrows could be colored as well. Any of the cyclic color maps discussed for the angles earlier, can be chosen. Figure 32 shows an example of this scenario. Not only does it make following the angle changes throughout the system easier, but it is also conducive to determining which areas swing together, for instance, rather quickly.

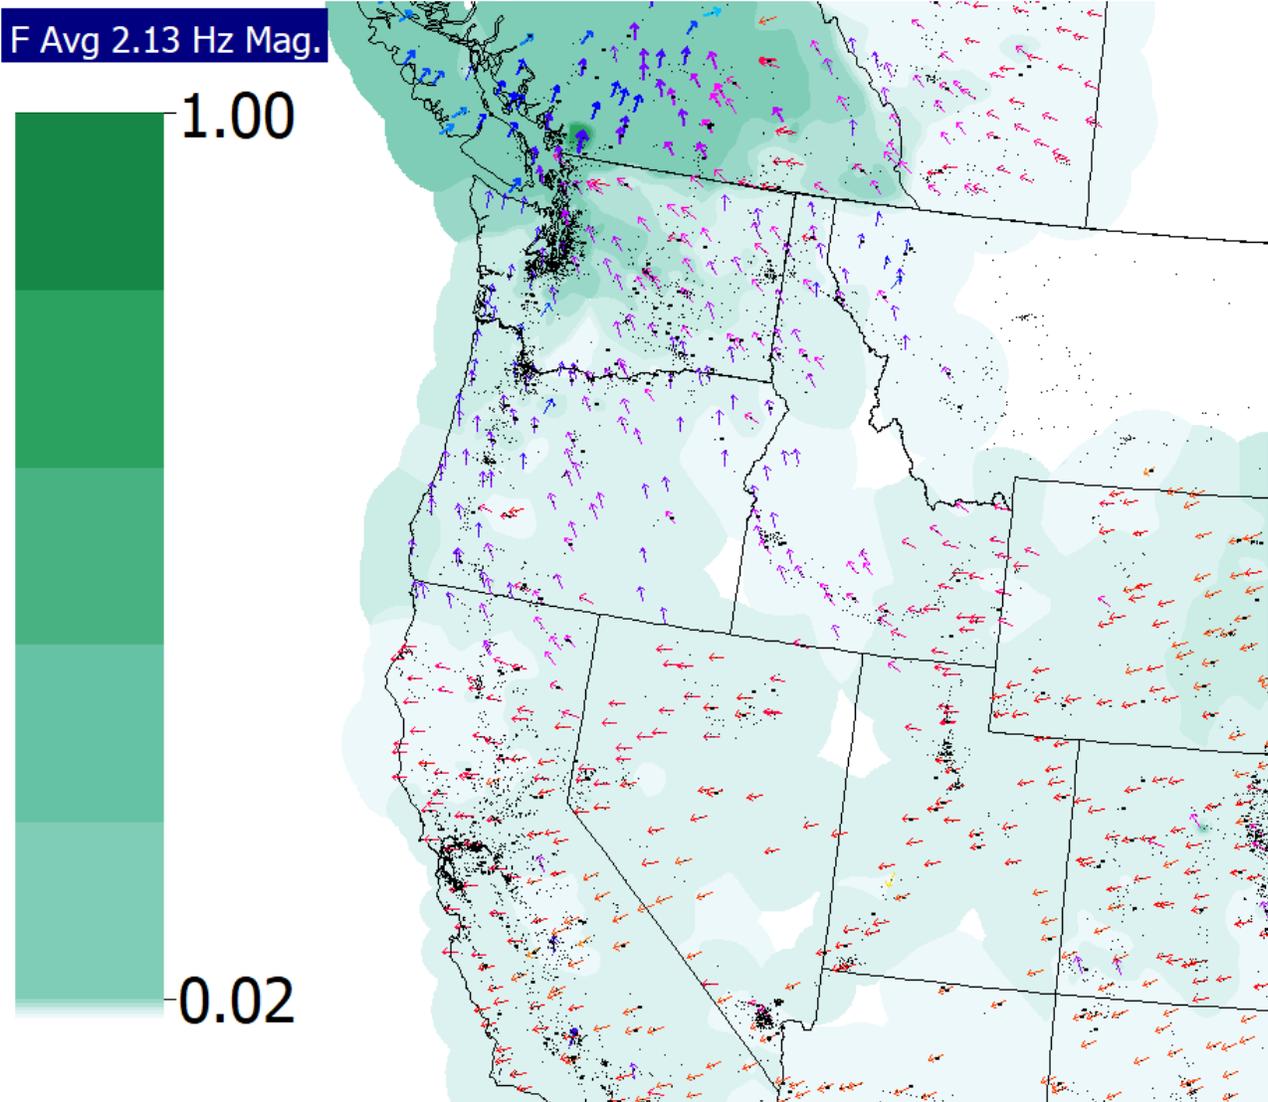


Figure 32. Visualizing 2.13 Hz mode details similar to Figure 31, except that colored arrows are used here

3.2 Voltages

Next, we perform modal analysis using the average of the voltage magnitudes of the buses at each substation as the input signals, with Table 2 showing the results. Following the procedure described earlier for frequencies, 20 iterations are eventually chosen during the optimal matrix pencil analysis with these signals. Since voltages are more localized in nature, it follows that there may tend to be more variations in the voltage response compared to frequencies, thus making it necessary to include more signals and iterations.

Table 2. Modal analysis results for Substation Average Voltage Magnitudes

	Frequency (Hz)	Damping (%)	Largest Weighted Percentage for Mode	Signal Name of Largest Weighted Percentage for Mode	Lambda
1	2.129	-0.036	89.6936	Substation Stave Falls V pu Average	0.0048
2	4.254	0.147	8.8966	Substation Stave Falls V pu Average	-0.0393
3	2.847	0.702	21.8311	Substation Seven Mile V pu Average	-0.1256
4	5.000	3.260	9.2921	Substation Hopkins Ridge V pu Average	-1.0248
5	0.970	3.439	82.5981	Substation Friant Hydro Facility V pu Average	-0.2096
6	0.909	3.851	54.0064	Substation Milford Sub. V pu Average	-0.2202
7	0.737	6.899	33.7547	Substation Riverview V pu Average	-0.3201
8	0.651	9.089	40.4142	Substation Clayton V pu Average	-0.3734
9	1.247	10.719	33.5581	Substation Honey Lake Power Plant V pu Average	-0.8447
10	0.848	11.490	29.9393	Substation Rowley V pu Average	-0.6164
11	1.127	14.608	32.8501	Substation Dexter V pu Average	-1.0454
12	0.044	15.695	35.2639	Substation Endwind V pu Average	-0.0443
13	0.358	16.600	46.0984	Substation Endwind V pu Average	-0.3784
15	2.176	18.887	24.5604	Substation Stave Falls V pu Average	-2.6302
16	4.254	20.072	29.0091	Substation Hopkins Ridge V pu Average	-5.4770
17	2.926	21.881	18.0170	Substation Stave Falls V pu Average	-4.1227
18	0.192	39.220	20.0309	Substation Anaconda V pu Average	-0.5149
19	0.087	97.375	94.6495	Substation South Meager V pu Average	-2.3284

As with the frequencies, we again visualize some of the modes, one at a time. Figure 33 shows the angles of the 2.13 Hz mode, but this time the input signals used are substation average voltages and not frequencies. Comparing with its frequency counterpart in Figure 29, one can immediately notice the several local variations in these angles. This is, again, due to the localized nature of reactive power and voltage. However, despite the local variations, the angles across most of the BPA footprint appear to be uniform. This can make oscillation control (if desired) over this footprint more manageable. Note that majority of the discontinuities and variations appear in the southern part of the system, while for the frequency results most of the southern region was oscillating together.

Given the similar values of damping, and the fact that the largest weighted percentage is again 90% with the signal being Stave Falls voltage, the source of the both the frequency and voltage 2.13 Hz oscillations is very likely the same. However, this may not always be true; e.g. the same cannot be claimed with as much confidence for the 1.12-1.13 Hz mode appearing both in the frequency and voltage signals (Table 1 and 2). They have different damping values and dominant substations.

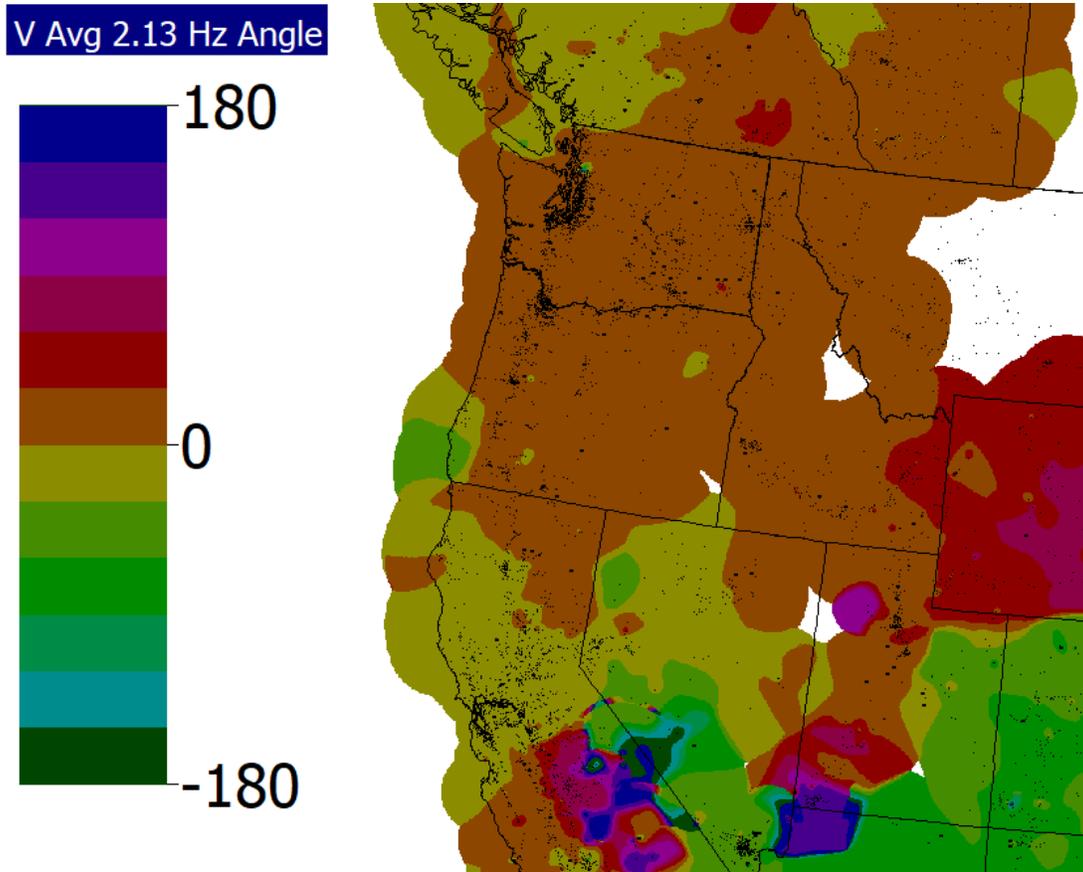


Figure 33. 2.13 Hz mode angles for voltage signals

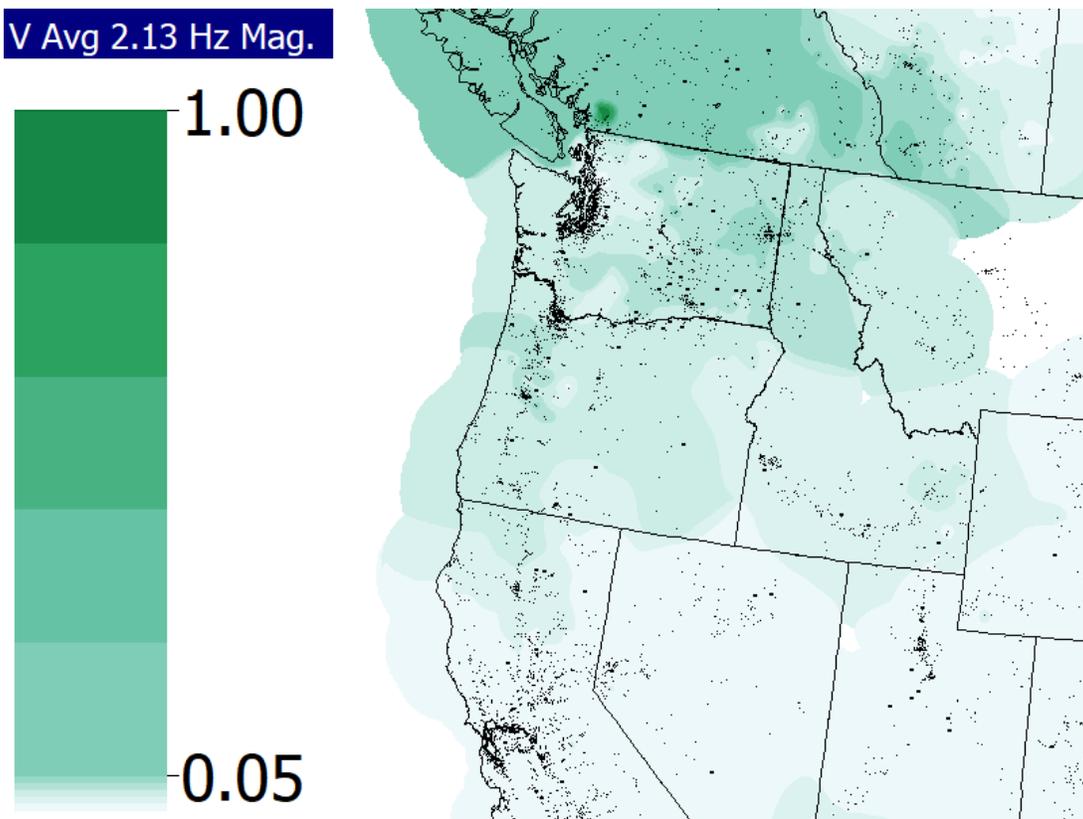


Figure 34. 2.13 Hz mode magnitude for voltage signals

Next, we inspect the 0.91 Hz mode from Table 2. Its details are visualized in Figure 35. Here we see a lot more local variation in the angles aided by the colored arrows, even in the BPA footprint. In fact in OR, a bit inland from the coast, substations in nearby locations are oscillating almost against each other, indicated by the red and cyan arrows. Overall, the variations indicate that there are some interesting dynamics associated with this mode in the BPA footprint.

The magnitude contour for this mode points to a region in UT. Substation Milford, which has the largest weighted percentage for this mode (54%) actually lies in CA. However, for the 0.91 Hz mode in the frequency signals, the corresponding “largest weighted percentage of the mode” signal is from Substation Blundell (93%). This substation is located in the same region where the dominant magnitudes are found with voltage signals in Figure 35. Thus, there is a strong coupling between the voltage and frequency oscillations at this frequency. Also note that Blundell is a generating substation, whereas Milford does not contain generators. That may be a contributing factor to the difference in the percentages.

Another example of a common mode between the voltage and frequency signals is the 0.97 Hz mode. From the frequencies, the dominant substation is Helm (72%) whereas among the voltage signals Substation Friant Hydro (82.6 %) is the largest weighted percentage signal. In this case the latter is a generating substation, while Helm is not. From several example studies so far, typically substations with largest percentages more than around 85% happen to be contain generators. Moreover, substations with such large percentage values are also likely to be sources of local plant/generator or forced oscillations. The last statement was already corroborated with the 2.13 Hz mode example, both in the frequency and voltage analyses.

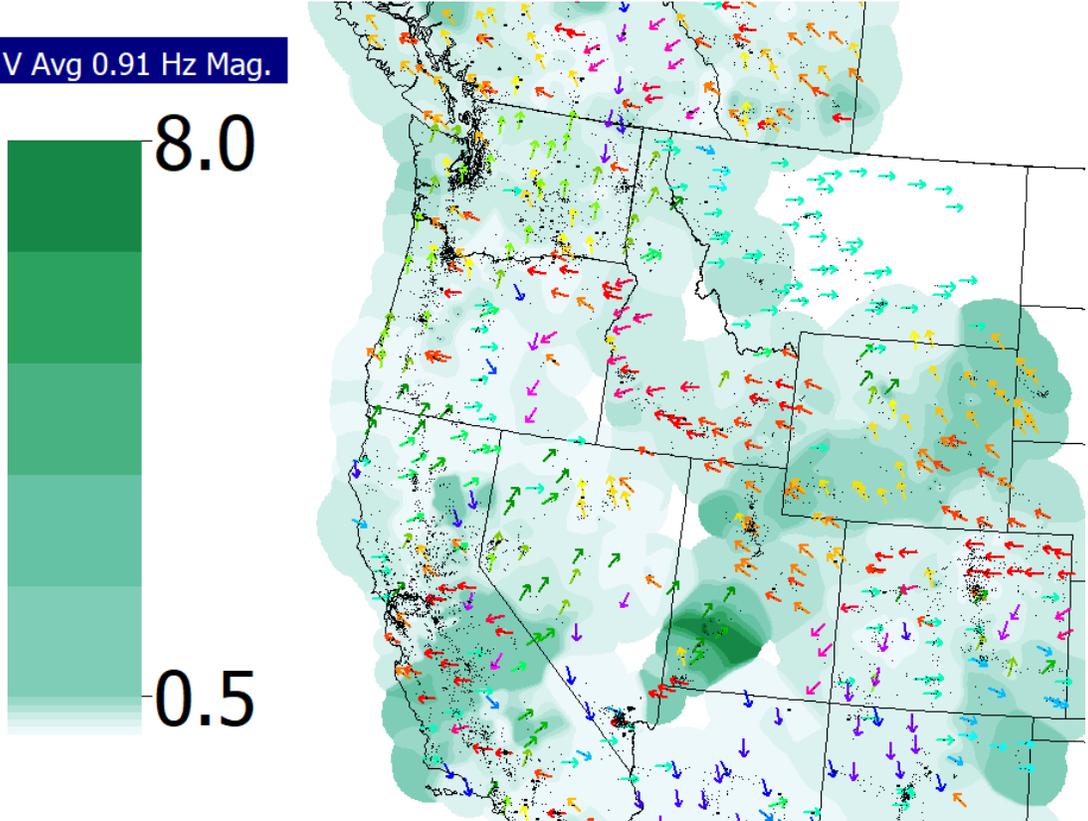


Figure 35. 0.91 Hz mode magnitude and angles using colored arrow GDVs. The largest percentage signal lies in CA, but many of the dominant signals lie in UT.

We continue with the 2.85 Hz mode, followed by a discussion of the 1.25 Hz mode as well. The 1.25 Hz is much better damped than the 2.85 one. This is clear from Figure 37 and Figure 38. However there is one

more important aspect to note in these figure, which is the Y axis scale. Since it represents the deviation in the frequency values at various substations from 60 Hz, it shows that the 1.25 Hz mode causes a much bigger impact in terms of the deviation, although the oscillation dies out quickly. For example, the ratio of the first peaks of the deviations is $0.12/0.0028 = 43$. To make this impact more evident, a metric could be added to the modal analysis results summary to quantify the deviations caused by each mode (minimum, maximum, mean values). In the examples discussed so far, there has been atleast an order of magnitude difference between the deviations (Y axes of Figure 28, Figure 37, and Figure 38) caused by the modes.

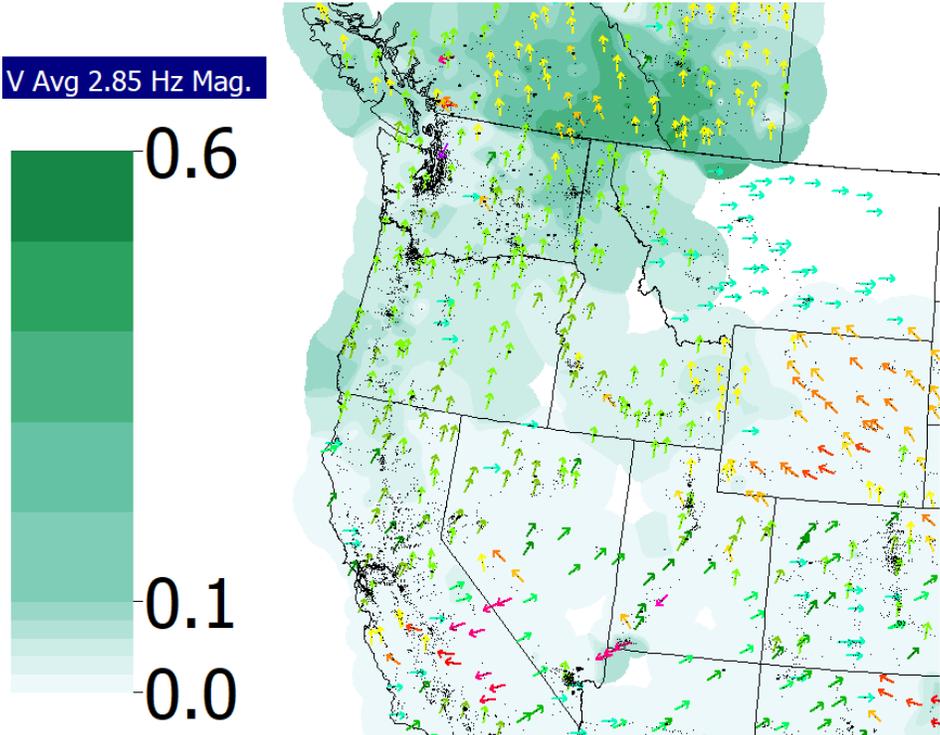


Figure 36. 2.85 Hz mode magnitude and angles

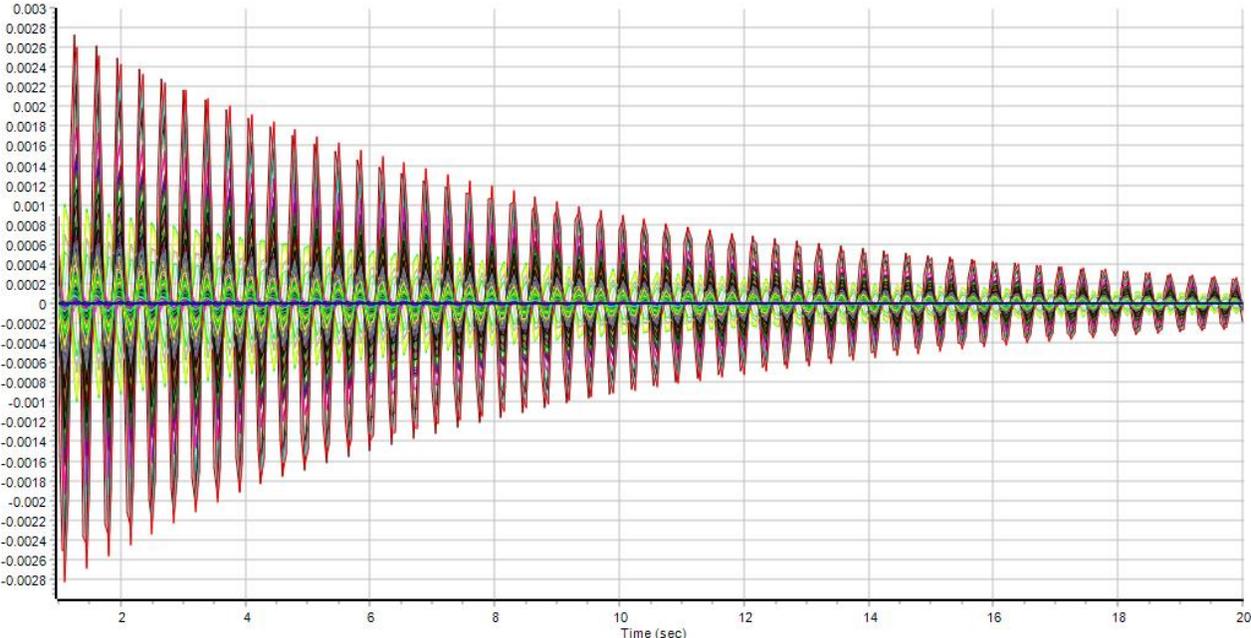


Figure 37. 2.85 Hz mode components

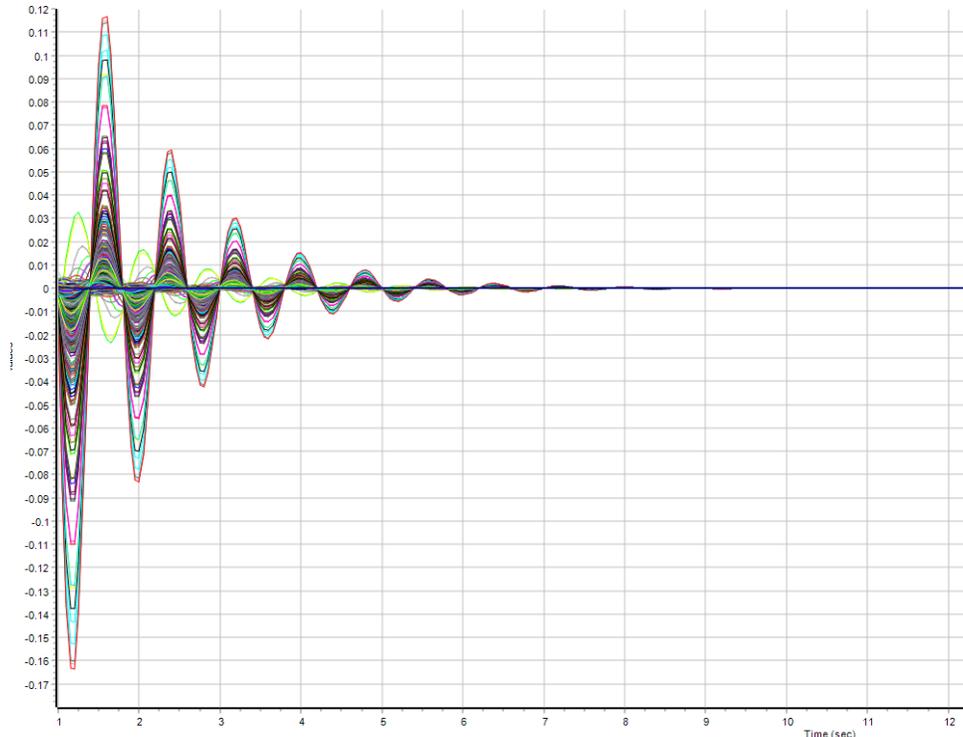


Figure 38. 1.25 Hz mode components. Note the Y axis i.e. the frequency deviations.

We conclude this section with a visualization of the 1.25 Hz mode. The reason for multiple examples in this subsection is to illustrate how the angles do not have a fixed pattern, or in other words, angles associated with voltage signals may not always be localized or have several local variations. In fact, in case of the 1.25 Hz voltage mode, the angle contour of Figure 40 almost resembles one of the lower frequency WECC modes, in that the angles have the same value across a wide area, in this case most of WECC which is then anti-phase with NM, AZ and parts of two other states. Alberta has some interesting angle discontinuities.

Another unique attribute apart from the uniformity in angles is that unlike previous voltage examples, there is no corresponding frequency signal 1.25 Hz mode with similar damping (though 1.12 Hz with 9.8% damping comes close). More analysis is needed to confirm whether these two oscillations originate from the same source. The largest weighted percentage signals are located nowhere close to each other unlike previous situations (CA vs WA).

There are other general observations about calculated modes that need further consideration. For example, at times they appear in multiples of 2, with the same dominant signal. This is seen in Table 2 (2.129 and 4.254 Hz at Stave Falls) as well as Table 1 (0.726 and 1.404 Hz at Gorge SCL). It could be a property of the system, or an error. Care should be taken with factors such as sampling time, maximum frequency etc. to ensure aliasing does not occur.

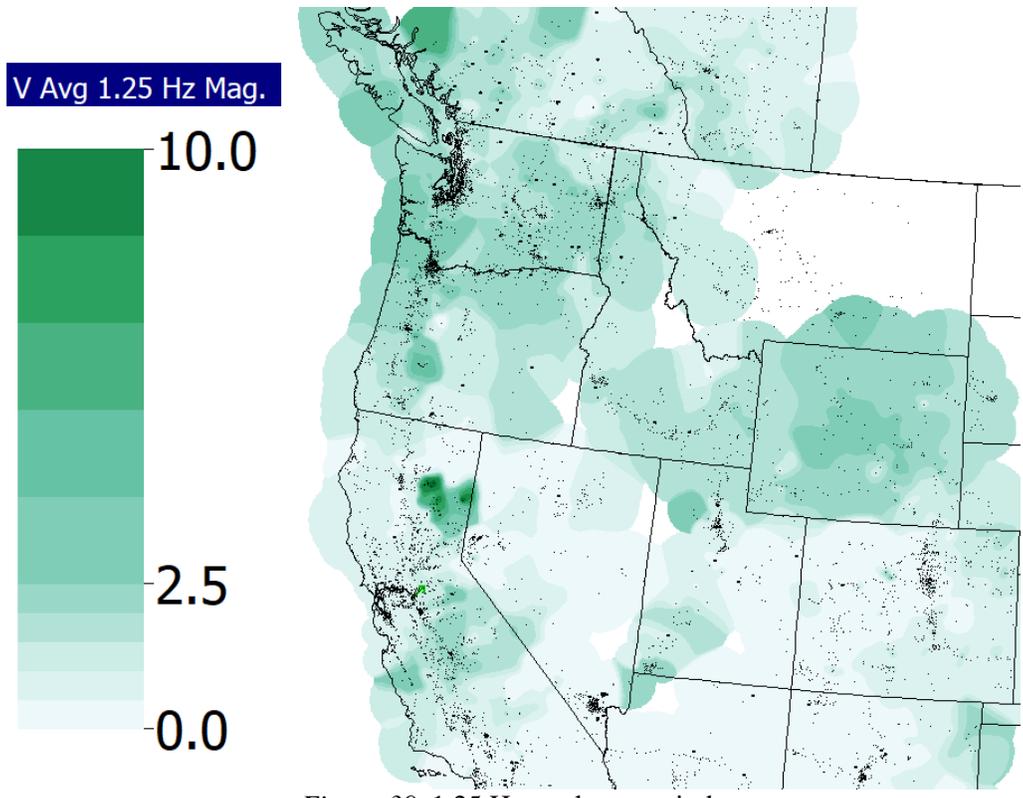


Figure 39. 1.25 Hz mode magntiude

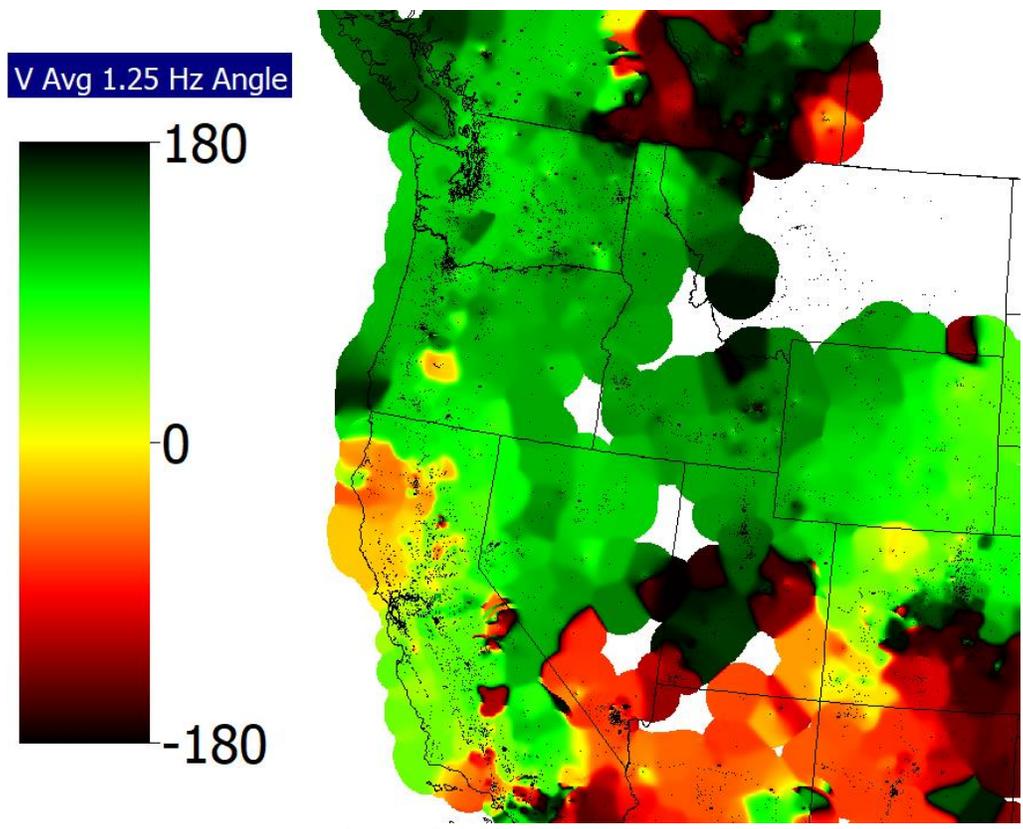


Figure 40. 1.25 Hz mode angles

3.3 Effect of Reducing Number of Input Signals – PMU Deployment Constraints

So far, we have assumed that simulation results or measurements are available from all substations in the system. In reality, however, there are a limited number of PMUs deployed which measure data from important substations. These are typically, 1) high voltage, transmission-level substations with maximum nominal voltage levels of 230 kV and higher, and 2) substations containing generating units. To assess the impact of this practical concern, we stored results only at these substations, with an additional constraint that the generating substations recorded should have a maximum capacity of 100 MW or higher. This helped further narrow the signal space to a total of 1503 substations, which is about 1/5th of the number of signals in the original simulation.

We use substation average frequencies, and the 0.36 Hz system-wide mode as the example. Figure 41 shows the magnitude by the size and thickness of the arrows, and angle by the direction and color. Contours may not work well in this case since the loss of geographic data points may cause inaccuracies due to the inherent interpolation. Since GDV points are discrete and independent of each other, these can be used effectively. Comparing this with the results obtained using all the signals in Figure 20 and Figure 21, the similarities in the angle and magnitude distribution can be observed. For instance, Alberta and southern WECC oscillate together, against rest of WECC, particularly Pacific Northwest. Magnitude-wise, the northern and southern ends of the network are strongest, with weaker magnitudes in the middle along NV, UT, and WY.

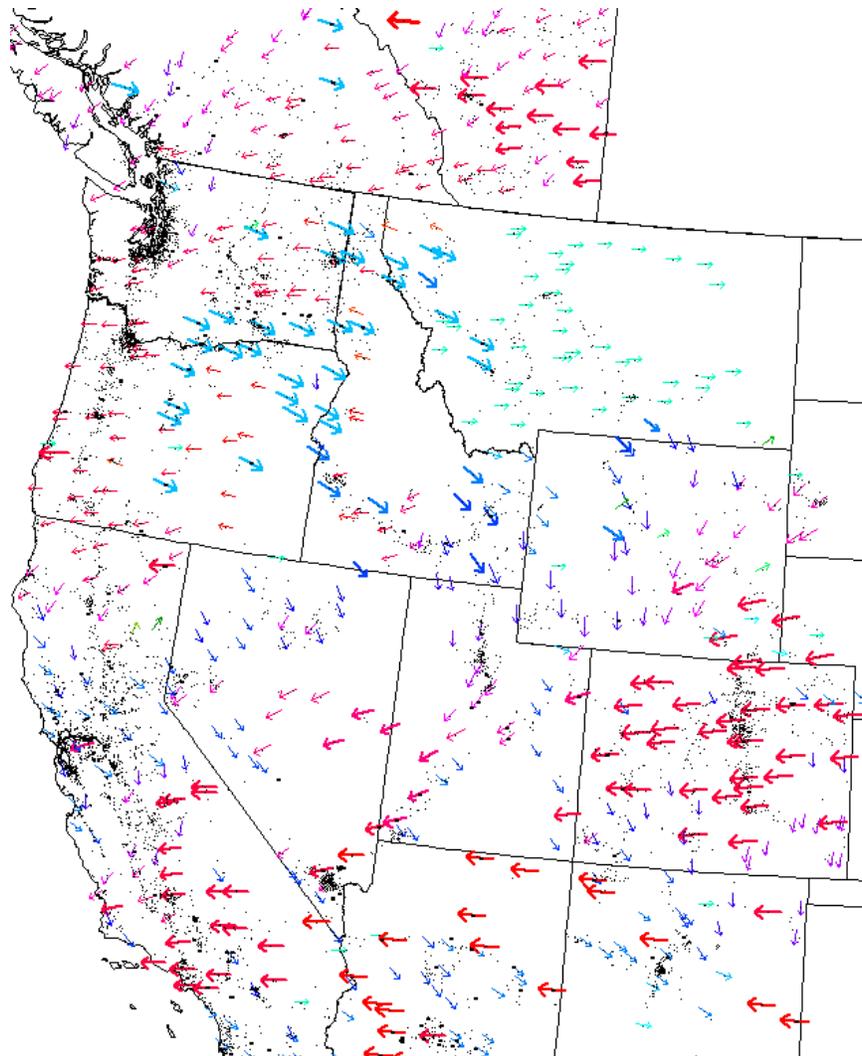


Figure 41. GDV using arrows for the 0.36 Hz mode, using only 1503 substation average frequencies

As far as the overall modal analysis results are concerned, we again use the same conditions as before (time window, maximum frequency, number of iterations). The solution ends up with the highest cost slightly less than the case with all signals. The results of the modes themselves are quite similar, especially with regards to the important ones discussed in the report. Slight discrepancies in frequency and damping values do exist; here they do not have much of an impact. Sensitivity studies with varying number and locations of signals can help determine impacts of these parameters more systematically.

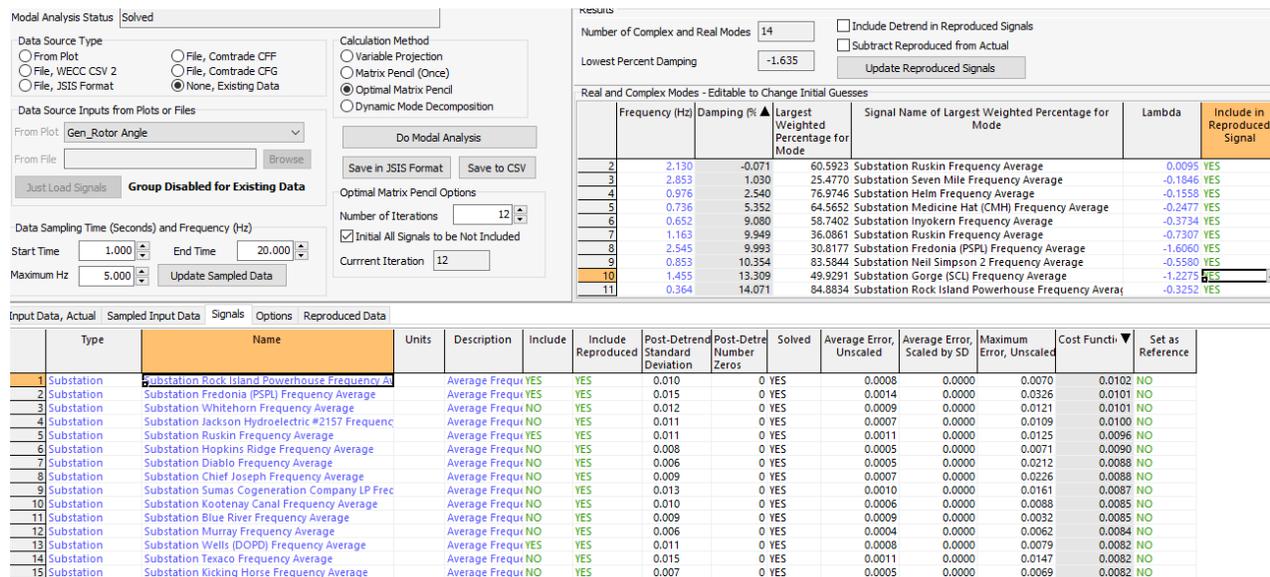


Figure 42. Modal analysis results for 1503 substation average frequencies

The previous 0.36 Hz mode example was that of a major system-wide mode. For the local mode of 2.13 Hz, Figure 43 shows the angles using the subset of signals. The results do not resemble the original simulation (i.e. Figure 29 or Figure 32), as in the case of the 0.36 Hz mode.

In the southern part of WECC especially where angles are supposed to be uniform, several variations are observed, some of them anti-parallel in close proximity (orange and dark blue). This could be attributed to the fact that this is a local oscillation so the effects are localized and experienced near the source which is far in the northwest. This in addition to the absence of enough number of signals carrying information about the impact of the 2.13 Hz mode away from the source could be likely causing the errors in the angle calculations in southern WECC. Note that “errors” here is with reference to the solution with all 7300 signals included.

The local angle estimates are still reasonable; there is a consistent phase shift between BPA and Alberta, with a slight exception in BC, across the two scenarios.

To summarize, this approach looks promising even when the number of input signals is cut down by a factor of 5. The estimation of the frequencies and damping ratios of the key modes is not significantly affected, nor does the optimal matrix pencil method suffer much in terms of the worst cost function. At most it may pose some challenges in visualization (with alternatives already available), and in analyzing local modes from a wide-area perspective. Future research can focus on these answering these specific questions to reap the full benefits of this technique.

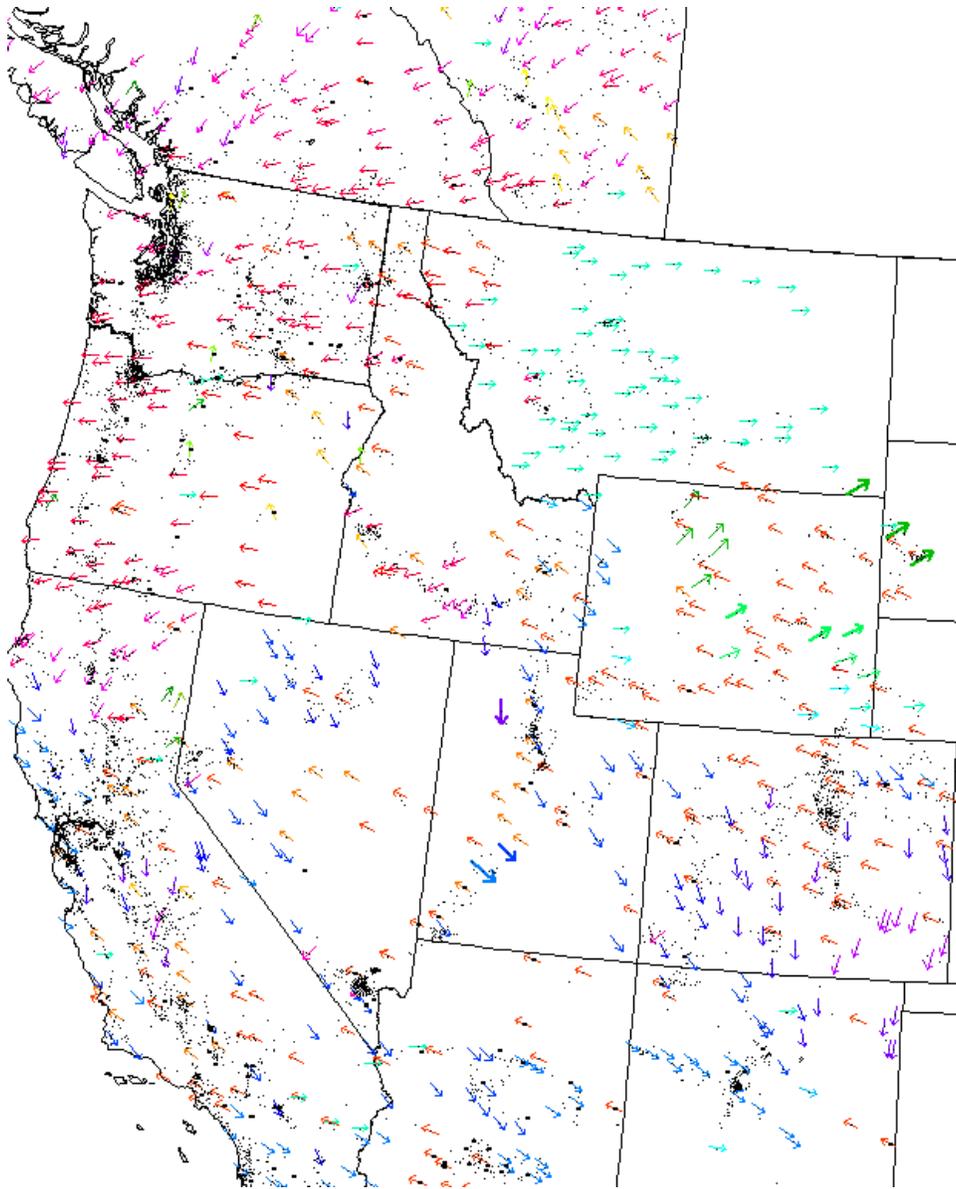


Figure 43. 2.13 Hz mode details calculated using 1503 signals

In the next and final section, we discuss an analytical method to locate the source of oscillations, and apply it to the 2.13 Hz mode. However a different set of input signals is used, as required by this method.

4 Source of an Oscillation

Several techniques exist in the literature to determine sources of oscillations. Utilizing the previously mentioned visualization of particular modes, the approach of [12] can be simplified to just show the flow of the power associated with the sustained oscillation frequency. As opposed to using a band-pass filter in [12], here we first do the modal analysis. Define the mode j component of the voltage angle at bus k as,

$$\theta_{k,j}(t) = A_{k,j} e^{\sigma_j t} \cos(\omega_j t + \phi_{k,j}) \quad (7)$$

For visualizing the source of sustained oscillations assume $\sigma_j = 0$ and recognizing that the source of the oscillation will have a leading phase angle, shift all the $\phi_{k,j}$ so it is zero at the bus with the largest $A_{k,j}$ and set $t = 0$. Then, $\theta_{k,j}(t) = A_{k,j} \cos(\phi_{k,j})$.

Then we approximate the mode j power flow on all lines between buses m and n as

$$P_{mn} \approx \frac{1}{X_{mn}} (A_{k,m} \cos(\phi_{k,m}) - A_{k,n} \cos(\phi_{k,n})) \quad (8)$$

The source of the sustained oscillations is the net source of the mode j power flow. Note that this method is only applicable to bus voltage angles.

Revisiting the previous example, we now perform modal analysis on the bus voltage angles (not substation quantities like before). The analysis again detects the negatively damped 2.13 Hz mode. The goal of this example is to accurately determine the source of this oscillation based on the above theory. Note that a likely source was discussed in the previous section, using the magnitudes of the substation frequency and voltage oscillations for this mode.

Type	Name	Units	Description	Post-Detrend Standard Deviation	Angle (Deg)	Magnitude Unscaled	Magnitude Scaled by SD	Cost Function
1 Bus	Bus SFN 13G1 V angle		Angle (deg)	1.758	-15.216	1.110484	0.632	0.0135
2 Bus	Bus SFN 60 V angle		Angle (deg)	1.487	-20.141	0.232493	0.156	0.0050
3 Bus	Bus SLK 2 V angle		Angle (deg)	1.487	-24.476	0.229383	0.154	0.0050
4 Bus	Bus SLK 60 V angle		Angle (deg)	1.487	-24.232	0.222547	0.150	0.0048
5 Bus	Bus ALU 60 V angle		Angle (deg)	1.488	-29.194	0.212987	0.143	0.0046
6 Bus	Bus ALU 6G1 V angle		Angle (deg)	1.501	-77.657	0.190610	0.127	0.0041
7 Bus	Bus HNY 25 V angle		Angle (deg)	1.472	-21.757	0.186896	0.127	0.0051
8 Bus	Bus HNY 60T2 V angle		Angle (deg)	1.473	-21.272	0.178696	0.121	0.0048
9 Bus	Bus HNY 60T1 V angle		Angle (deg)	1.473	-21.319	0.176320	0.120	0.0048
10 Bus	Bus MRG 25A V angle		Angle (deg)	1.467	-22.160	0.159631	0.109	0.0049
11 Bus	Bus MRG 25B V angle		Angle (deg)	1.467	-22.160	0.159629	0.109	0.0049
12 Bus	Bus MRG25T12 V angle		Angle (deg)	1.467	-22.160	0.159627	0.109	0.0049
13 Bus	Bus MRG60T12 V angle		Angle (deg)	1.468	-21.688	0.152773	0.104	0.0046
14 Bus	Bus MRG60T11 V angle		Angle (deg)	1.468	-21.684	0.151649	0.103	0.0046
15 Bus	Bus MRG60T13 V angle		Angle (deg)	1.468	-21.555	0.148013	0.101	0.0045
16 Bus	Bus RUS 60 V angle		Angle (deg)	1.486	-21.849	0.148497	0.100	0.0042
17 Bus	Bus RUS 13G3 V angle		Angle (deg)	1.486	-21.849	0.148497	0.100	0.0042
18 Bus	Bus MIS 12P V angle		Angle (deg)	1.478	-22.215	0.139719	0.095	0.0041
19 Bus	Bus MIS 60T2 V angle		Angle (deg)	1.478	-22.367	0.139328	0.094	0.0041
20 Bus	Bus SKW 4G1 V angle		Angle (deg)	1.485	-50.352	0.130469	0.088	0.0040
21 Bus	Bus MIS 25 V angle		Angle (deg)	1.477	-25.473	0.128923	0.087	0.0041
22 Bus	Bus SKW 25 V angle		Angle (deg)	1.477	-25.497	0.128912	0.087	0.0041
23 Bus	Bus SVD 25 V angle		Angle (deg)	1.479	-23.137	0.126189	0.085	0.0040
24 Bus	Bus SVD 60 V angle		Angle (deg)	1.480	-22.941	0.123578	0.084	0.0040
25 Bus	Bus PKL 25 V angle		Angle (deg)	1.465	-23.972	0.117598	0.080	0.0045
26 Bus	Bus MIS 60T1 V angle		Angle (deg)	1.478	-23.533	0.113175	0.077	0.0039
27 Bus	Bus MIS 14P V angle		Angle (deg)	1.478	-23.419	0.112444	0.076	0.0038
28 Bus	Bus PKL 60T2 V angle		Angle (deg)	1.469	-23.390	0.110836	0.075	0.0041
29 Bus	Bus PKL 60T1 V angle		Angle (deg)	1.469	-23.379	0.110052	0.075	0.0041
30 Bus	Bus BAL 8P V angle		Angle (deg)	1.470	-23.144	0.108234	0.074	0.0040

Figure 44. 2.13 Hz mode detected by modal analysis of the bus voltage angles

We make use of GDVs to visualize these flows, belonging to a particular mode, in or out of each substation. The sizes and colors of the GDVs can be set based on the magnitude and direction of the power flow. Figure 45 shows the use of ovals to represent the key substations. Red ovals indicate the mode power sources, whereas yellow denote sinks. Due to the bus fault and opened lines, the fault location may appear to be a large source (indicated by a very large calculated power value compared to the “true” source), especially if it is close to the actual source of oscillation. To mask this false positive, we set a reasonable threshold for the power value and indicate the fault location by a green oval. The yellow ovals have been intentionally scaled up since they were being hidden by the red ovals. This may work for a prototype but in actual working environment, an alternative approach might be better. For example, transparent GDVs can be implemented to view closely located / almost overlapping sources and sinks.

This method detects the source to be in BC, just off the US border. On further inspection, it is Stave Falls substation itself, thus confirming our hypothesis from the previous sections.

While there typically may be one big source, the sinks can be small and several, and distributed across an area. In this example, several small yellow ovals are present in BC and WA. Figure 46 shows a zoomed in version, to better view these smaller ovals.

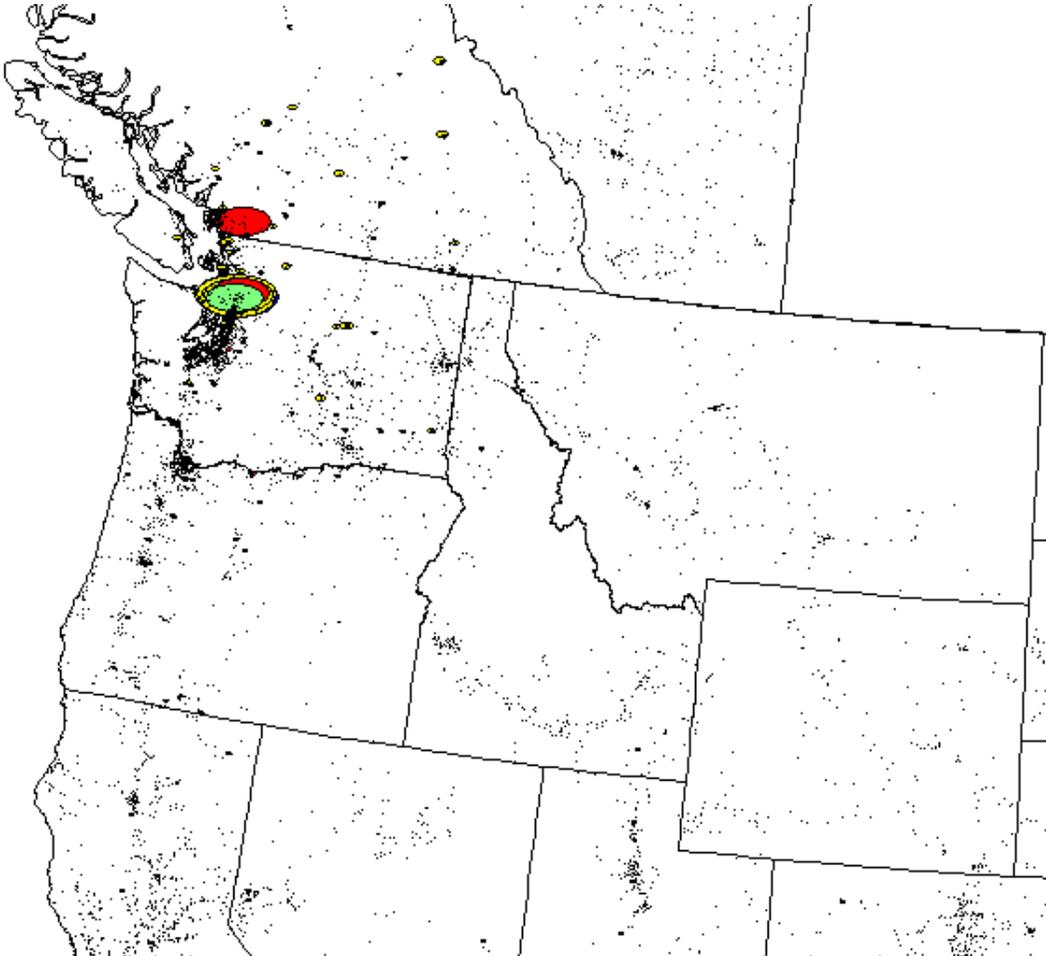


Figure 45. GDV with substations represented by ovals showing sources (red) and sinks (yellow) of the 2.13 Hz mode power flow. Size of ovals is proportional to the power. The green oval covering the red oval (which should not be considered a source) is the location of the fault. Note that the yellow oval is scaled up 1.5 times to prevent it from being hidden.

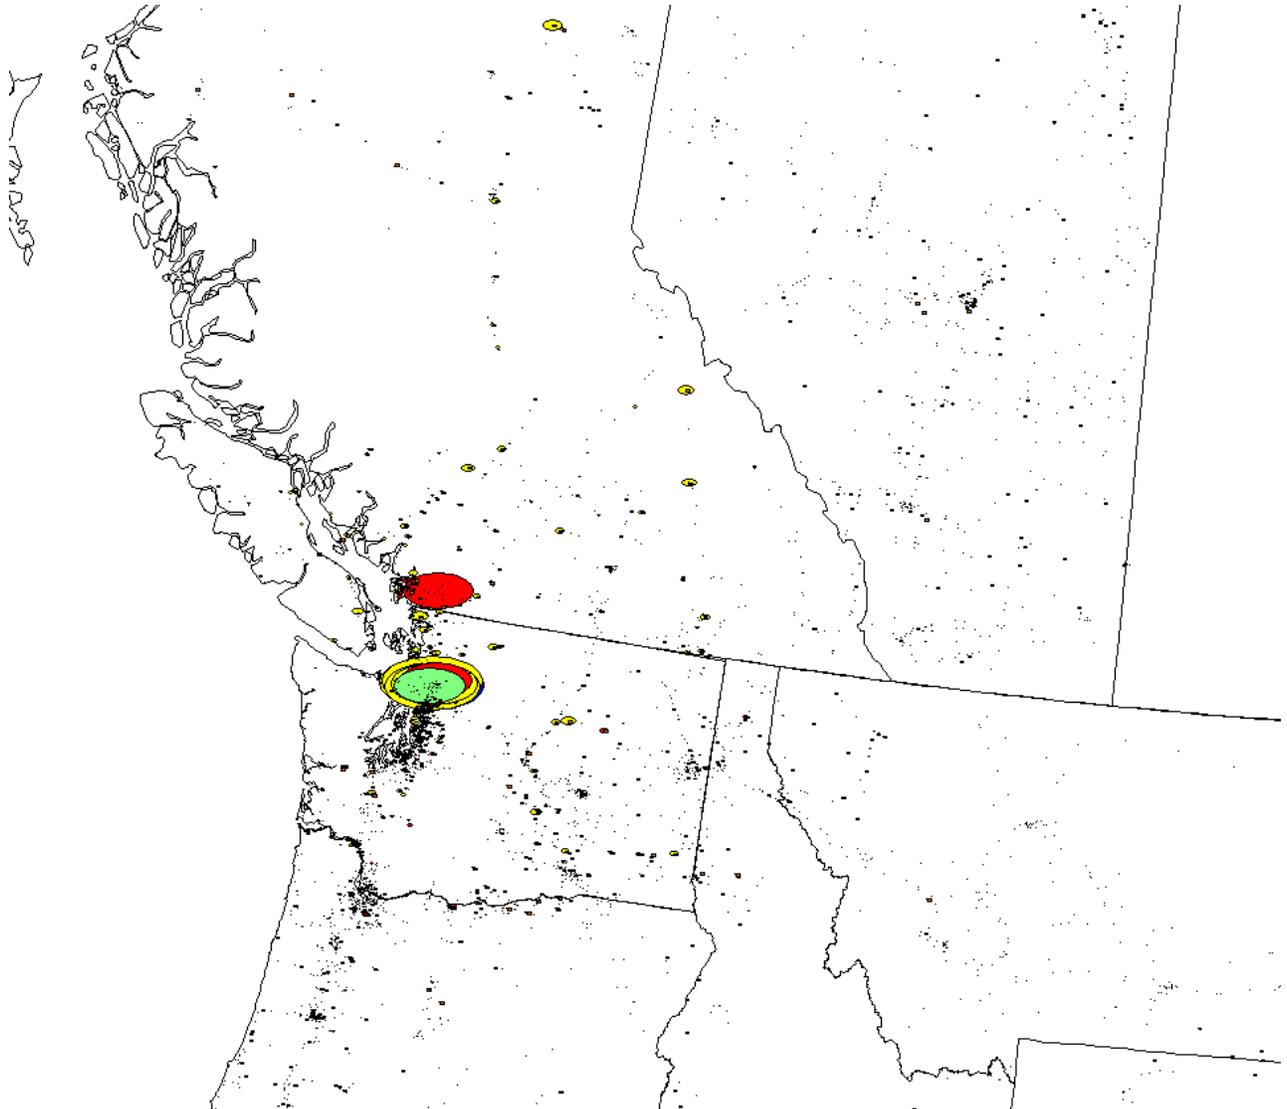


Figure 46. Zoomed in view of Figure 45. Notice the small yellow ovals in WA and BC which act as sinks for the oscillating power.

Last, we try to find the source of the 0.91 Hz oscillation since it was causing major deviations. We ran the same analysis and plotted the GDVs the power flow of that mode. Instead of one clear cut source, atleast 3-4 potential sources of comparable magnitude are observed from the GDV. Nonetheless, the partially overlapping large red ovals located near northwest WA can be deemed as the source of this flow. The sinks, which are again exaggerated to highlight them, are more pronounced and more prevalent in and around the BPA footprint. This could be attributed to the larger quantity of power associated with the larger frequency deviation, as opposed to the 2.13 Hz case.

Even if this method is not directly used or able to accurately detect the exact source of oscillation each time, it can help visualize the distribution of power associated with a mode. This can help in tuning dynamic model parameters of devices in the vicinities of these sources and sinks, and plan other measures for frequency control.

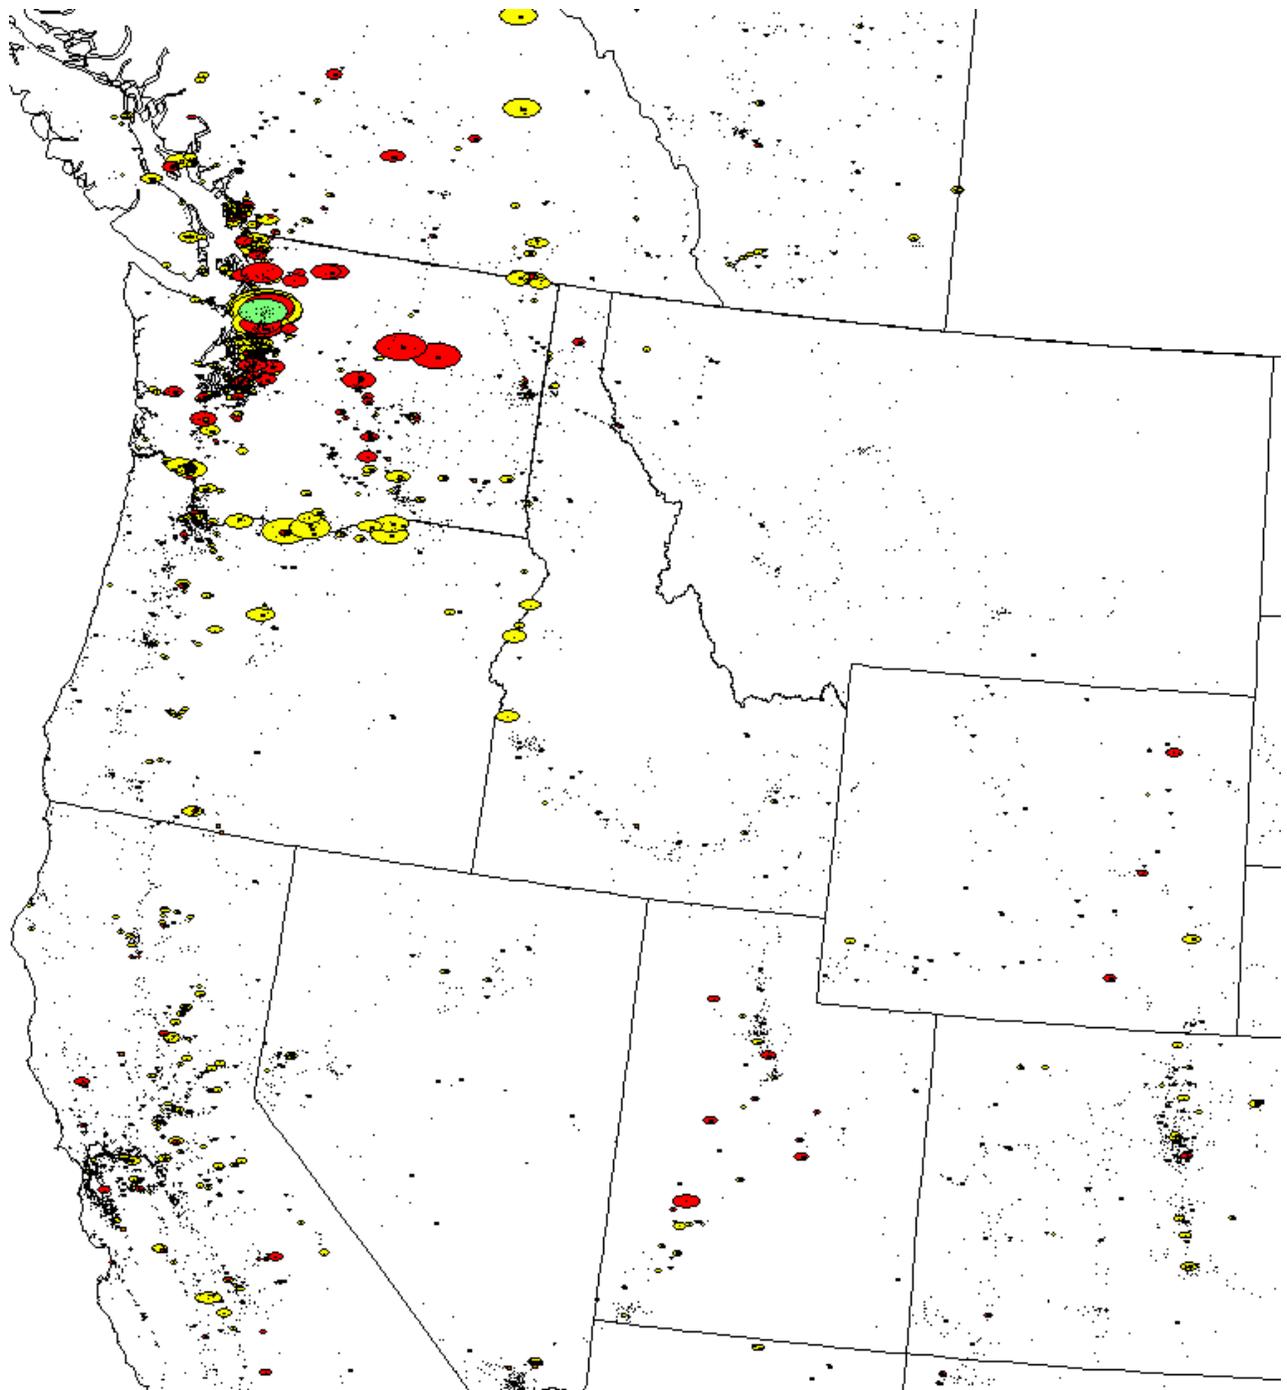


Figure 47. Source and sinks of the oscillating power for the 0.91 Hz mode

5 References

- [1] J. F. Hauer, "Application of Prony analysis to the determination of modal content and equivalent models for measured power system response," in *IEEE Transactions on Power Systems*, vol. 6, no. 3, pp. 1062-1068, Aug 1991.
- [2] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," in *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 5, pp. 814-824, May 1990.
- [3] L. L. Grant and M. L. Crow, "Comparison of Matrix Pencil and Prony methods for power system modal analysis of noisy signals," *2011 North American Power Symposium*, Boston, MA, 2011, pp. 1-7.
- [4] A. R. Borden and B. C. Lesieutre, "Variable Projection Method for Power System Modal Identification," in *IEEE Transactions on Power Systems*, vol. 29, no. 6, pp. 2613-2620, Nov. 2014.
- [5] A. B. Birchfield and T. J. Overbye, "Convergence characteristics of the variable projection method for mode extraction," *2017 IEEE Texas Power and Energy Conference (TPEC)*, College Station, TX, 2017, pp. 1-6.
- [6] "Modes of Inter-Area Power Oscillations in Western Interconnection," WECC JSIS Report, November 2013.
- [7] D. Kosterev, "Control Room Applications for Oscillation Detection," WECC JSIS Meeting, September 2015 [Online] Available: <https://www.wecc.biz/Administrative/18%20Kosterev%20BPA%20Oscillation%20Damping.pdf>
- [8] E.R. Tufte, *Envisioning information*. Chapter 3, Graphics Press, Cheshire, CN, February 2017.
- [9] C. Ware, *Information Visualization: perception for design*. chapter 6, 3rd ed. Elsevier, MA, 2013.
- [10] T.J. Overbye, E.M. Rantanen and S. Judd, "Electric power control center visualization using Geographic Data Views," *2007 iREP Symposium - Bulk Power System Dynamics and Control - VII. Revitalizing Operational Reliability*, Charleston, SC, 2007, pp. 1-8.
- [11] R. H. Byrne, R. T. Elliott, J. S. Neely, C. A. Silva-Monroy, D. A. Schoenwalk, and L. L. Grant (Sandia National Laboratories), "Renewable Source Controls for Grid Stability," *Sandia Report*, SAND2012-10416, December 2012.
- [12] S. Maslennikov, B. Wang, E. Litvinov, "Locating the Source of Sustained Oscillations by Using PMU measurements," *IEEE PES 2017 General Meeting*, Chicago, IL.