NASPI-2018-TR-001

Synchrophasor Monitoring for Distribution Systems:

Technical Foundations and Applications

A White Paper by the NASPI Distribution Task Team

January 2018

Editor:

Alexandra von Meier - UC Berkeley

Contributing Authors (in alphabetical order):

Reza Arghandeh - Florida State University

Kyle Brady - UC Berkeley

Merwin Brown – UC Berkeley

George R. Cotter – Isologic LLC

Deepjyoti Deka – Los Alamos National Laboratory

Hossein Hooshyar – Rennselaer Polytechnic Institute

Mahdi Jamei – Arizona State University

Harold Kirkham – Pacific Northwest National Laboratory

Alex McEachern – Power Standards Lab

Laura Mehrmanesh – UC Berkeley

Tom Rizy – Oak Ridge National Laboratory

Anna Scaglione – Arizona State University

Jerry Schuman – PingThings, Inc.

Younes Seyedi – Polytechnique Montreal

Alireza Shahvasari – UC Riverside

Alison Silverstein - NASPI

Emma Stewart – Lawrence Livermore National Laboratory

Luigi Vanfretti – Rensselaer Polytechnic Institute

Alexandra von Meier - UC Berkeley

Lingwei Zhan – Oak Ridge National Laboratory

Junbo Zhao – Virginia Tech

Contents

1.0 Introduction 5
1.1 Premise of Distribution PMUs 6
1.2 What's new? Synchrophasor technology 7
1.3 Why bother? High-value uses for distribution monitoring 8
1.4 Context for Distribution Sensing and Measurement 8
2.0 Brief Synchrophasor Definitions 10
3.0 Distribution-Specific Applications and Use Cases 11
3.1 Overview 11
3.2 Distribution Applications, Measurements and Data Needs 12
3.3 Event Detection, Classification and Cyber Attack Detection 15
3.4 Distribution Network Topology Detection 18
3.5 Model Validation and Phase Identification 19
3.6 Equipment Health Diagnostics 21
3.7 DG-Load Disaggregation 22
3.8 Fault Location 23
3.9 Network-Level Steady State Estimation 24
4.0 Synchrophasor networks and implementation 26
4.1 Distribution-level PMUs and monitors 26
4.2 Data quality and noise 27
4.3 Application Sensitivity to Data Quality 29
4.4 Transducer Errors 30
4.5 PMU Placement: Theory 33
4.6 PMU Placement: Practical Aspects 35
4.7 Communications and Data Collection 36
4.8 BTrDB 37
4.9 Cybersecurity and Synchrophasors 39
4.10 Time Synchronization Requirements 39
Appendix A: Synchrophasor Representation and Quantities 41
A.1 Phasor Representation 41
A.2 Total Vector Error 45

A.3 Comment on TVE 48
A.4 Phasor Quantities – A Closer Look 50
Appendix B: Public Distribution Synchrophasor Efforts and Deployments 55
References 58

1.0 Introduction

The purpose of this white paper is to provide an introductory reference for industry and academic practitioners interested in exploring the use of synchrophasors and other time-synchronized measurements for supporting power distribution system planning, operation, and research.

The mission of the North American Synchrophasor Initiative's Distribution Task Team (DisTT) is to foster the use and development of capabilities of networked phasor measurement units (PMUs) and other time-synchronized monitoring devices at the utility distribution level, beyond the substation. This group shares information in support of effective research, development and deployment of distribution PMUs and their applications. In doing so, it aims to create a community for solving technical and other challenges specific to distribution PMU technology and context.

These activities are motivated by the belief that effective measurement and analytics for the electric grid, including the distribution level, represent an important enabling technology for electric power quality, reliability, grid resilience, and sustainability – especially given the growing significance of diverse and renewable resources.

Significant changes are occurring at the periphery of the grid – more distributed generation and storage on customer premises, more customer-initiated demand response, electric vehicles and other changing customer load characteristics -- that necessitate better situational awareness and insight into distribution system conditions and performance. Although power distribution systems have not been highly instrumented in the past, there is growing interest in applying PMU-like technology at the distribution level.

PMUs take time-synchronized measurements of voltage, current and frequency that can tell grid operators what is happening, where, and when. Their precision in physical measurement, time resolution, and the ability to cross-reference locations affords a deeper insight into physical power flows than conventional sensor data. Operators and planners can use this information to manage the new level of variability, uncertainty, and opportunity in the modern grid.

This White Paper aims to summarize key issues in distribution synchrophasor technology, convey a sense of the broad spectrum of applications, and provide a foundation for defining and prioritizing the needs for continuing research and development in this area.

The paper first articulates the premise for the idea of deploying synchrophasor networks in distribution systems, while recognizing the inherent difficulties. Section 2 provides a glossary of basic synchrophasor terms as a quick introduction for readers unfamiliar with the technology. Section 3 highlights a set of selected applications and use cases for synchronized measurements in distribution systems. Section 4 then delves into some of the details of implementing synchrophasor networks such as device placement, data quality, and data management.

Appendix A offers a more rigorous introduction to the mathematics of synchrophasor quantities and measurement errors, intended as a foundation for future discussions of PMU device performance and data quality.

Appendix B lists distribution synchrophasor installations and research projects known to the authors at the time of this writing.

1.1 Premise of Distribution PMUs

Historically, electric grid planners and operators had limited information for understanding the status and behavior of the electric grid. Available information included measurements from supervisory control and data acquisition (SCADA) systems, typically available at several-second intervals from substations, and model data based on equipment ratings and specifications. The physical state variables of the a.c. network – specifically, the complex voltages, or time-shifted voltage waveforms at every node – were not directly observable, but could be estimated through these models. This solution worked well enough for many years. But given the growing uncertainties and complexities in grid planning and operations, these methods are increasingly becoming inadequate in time resolution, precision, accuracy and scope.

Transmission planners and operators were first to recognize the need for new tools that rely on advanced sensors and more comprehensive monitoring to better observe, understand and manage the grid. The challenge in transmission systems was comparing measurements across long distances (hundreds of miles) that would reveal physical interactions such as oscillations between generators, and be able to describe power flows and stability across an entire synchronous a.c. network. By comparison, distribution systems were simple and posed little need to observe their operation with much granularity in space or time.

But with the rapid growth in deployment of distributed energy resources, two-way electricity flows and new customer devices such as electric vehicles, there is a growing interest in sharper observation tools for the distribution grid. The possibility of new interactions among new and legacy devices, along with opportunities for more active and intelligent control, delivers value from measurements that are both precise and time-synchronized, making electrical events and responses observable and comparable between locations.

A high-value grid monitoring system will possess several characteristics:

- A high degree of time granularity, on the order of a sample per cycle, compared to current SCADA and EMS, which provide samples every few seconds;
- Fast communications access for real-time streaming of data for system recovery following disturbances;
- High-resolution data for off-line engineering analysis, and preferably in near real-time to enable operation support analytics;
- Deployment of a large number of measurement devices across the system, which implies both low-cost devices and easy installation;
- Precise time synchronization of measured data to enable comparison across many electrical locations on the grid;
- Data quality, availability and volume that are appropriate to serve the high-priority uses and monitoring needs of operational and planning tools.

1.2 What's new? Synchrophasor technology

Synchrophasor technology, using phasor measurement units (PMUs) offers all of the above characteristics. An individual PMU measures grid conditions at high speeds at a specific location and gives each measurement a precise time-stamp; a network of PMUs sharing data through a phasor data concentrator (PDC) or other aggregating platform enables the calculation of time-synchronized voltage and/or current magnitudes and phase angles (known as synchrophasors). There are substantial differences between PMU instruments suitable for use in transmission and in distribution systems, but their strategic significance is similar in that they both offer unprecedented situational awareness.

Precise synchronization makes it possible to directly observe the time shift in the voltage or current waveform as measured at different locations, which is a small fraction of a 60-Hz cycle. The voltage phase angle difference δ is a state variable, along with voltage magnitude, that can be thought of as physically driving a.c. power across a circuit or network. The ability to measure phase angle changes at multiple locations allows grid operators to detect and characterize grid behaviors such as oscillations that could not be observed by traditional lower-speed state estimation and monitoring systems, or even by dedicated, single-location modern high resolution measurements such as power quality monitoring. Until the advent of synchrophasor networks, grid situational awareness was limited by the relatively low precision of modeling estimates, the strictly local character of higher resolution measurements, and the lack of a precise time stamp on measurement data.

Once a PMU network is deployed for the purpose of detecting specific grid behaviors, the same network can be used to monitor traditional quantities, perhaps at greater resolution, and at little additional cost. Comprehensive real-time monitoring enables grid operators to observe and react to any significant change in the grid, such as electric power events at generators or customer facilities, and to archive those events for later analysis (as for forensic purposes or pattern identification) and training. Once the PMU archive has been created, users can apply a variety of analytic tools on those data, discovering new insights from old data as analytic capabilities evolve.

Synchrophasor capability may be built into devices such as relays or digital fault recorders (DFRs), and many PMUs that could be considered within the distribution realm are embedded in digital relays at distribution substations. The most prominent PMU used beyond the substation level today is called the micro-PMU (μPMU). In the future, PMU and networking capability may be embedded in devices such as inverters, transformers, etc.

Traditional electric monitoring systems used dedicated sensors to monitor specific grid conditions at specific locations for specific applications or uses. In contrast, the synchrophasor *network* monitoring approach collects a wide set of compatible grid condition data from many locations, for a broad and evolving set of uses. When only a small number of locations need monitoring, it is probably less expensive to use functionally-dedicated measurement tools in a few precise locations. But as there are many event locations of interest, and the use cases may expand over time, it is probably more cost-effective to use a PMU-like data collection system that can support multiple applications. If a PMU network is warranted for managing unknowns of high uncertainty, especially as the penetration of distributed resources increases, that crossover should come relatively quickly.

1.3 Why bother? High-value uses for distribution monitoring

A substantial fraction of electric utilities' financial and capital assets are spent on distribution networks. Despite good design, installation, operation and maintenance efforts, well over 90% of customers' electric outages occur due to problems occurring on the distribution system (rather than from transmission- or generation-level problems). Yet many North American utilities have limited amounts of monitoring on their distribution systems. Typically, there are some SCADA devices on subtransmission elements, and a growing number of advanced meter deployments that provide 15-minute power or energy readings (although often with delayed data delivery).

It is difficult to know what is happening on the distribution system without monitoring and measuring distribution-level and grid-edge activity. Distribution system managers could use high-quality, high-speed, wide-scale distribution-level monitoring – as feasible from a distribution-tailored synchrophasor network – planners and operators could use the data and appropriate analytical tools for many purposes, including:

- State or condition monitoring of the distribution system;
- Monitoring and analysis of customer-owned, behind-the-meter distributed generation and energy storage devices, enabling better forecasting and integration of those devices;
- Measurement and verification of customers' energy efficiency, demand response and load management activities (subject to appropriate privacy protections);
- Monitoring and analysis of significant end-user loads (for example, clusters of electric vehicle chargers);
- Identification of asset and equipment problems, including detection and advance warning of equipment operational issues and failures;
- Fault detection (including high-impedance faults), location and event forensics;
- Anomaly detection, including potential cyber-intrusions;
- Detection of previously unknown dynamic events (for example, control instabilities or oscillations) that are not recognizable with traditional monitoring.

Many or all of these capabilities can be used to cost-effectively improve distribution system design and operation and ultimately improve delivered reliability.

1.4 Context for Distribution Sensing and Measurement

Traditionally, grid operators and planners have had little in the way of explicit measurement information from primary and secondary distribution systems, between the substation and customer meters. They could largely get by with designing the system for the most extreme conditions – specifically, peak loads and faults – and then try to ensure that the grid operated within that expected range. The combination of radial topology, strictly one-way power flow out from the substation, and careful forecasts of customer load gave grid managers a reasonably good idea about the operating state of distribution circuits – i.e., voltages and currents between substation and customers – even without the benefit of empirical, real-time sensor data. Smart meters have added better information on customers' load demands and energy use, but meter data are limited in terms of the variables recorded (e.g., they may not include voltage), time resolution (most of the advanced meters presently deployed measure usage every 15 minutes), and latency (most systems do not collect smart meter data in real time).

The rapid growth of distributed energy resources (especially rooftop photovoltaic panels, but also other generation and energy storage) and high-demand loads (such as electric vehicles) has rendered the grid's traditional design basis invalid. Specifically, we may no longer assume one-way power flow, monotonic voltage profiles along distribution feeders, or net load currents that conform to a predictable pattern. For instance, customer load may be offset by solar generation behind the meter, and suddenly appear to spike when clouds pass, or when the PV inverter trips offline (as required) in response to a grid disturbance. In another example, batteries interconnected at the distribution level might participate in programs that aggregate output from multiple DERs for the purpose of providing transmission-level ancillary services (such as frequency regulation), and may charge and discharge without particular regard to their impact on the local distribution network. Since the distribution grid was not originally designed for two-way flow and interactions among diverse resources, there is a growing need to monitor the real-time performance, health and safety of distribution systems.

Distribution-level PMUs (such as µPMUs) permit the direct and precise observation of the state variables for a.c. power flow, synchronizing those measurements across many different locations on the grid.

This paper is tailored principally for the use of synchrophasor and other time-synchronized measurement systems on medium-voltage electric distribution systems. Section 2 below offers definitions of the measurement elements that a distribution-level synchrophasor system monitors (with additional technical detail included in Appendix A). Section 3 reviews the current known uses for distribution-level PMU monitoring. Section 4 reviews a number of technical considerations for distribution-level synchrophasor system design and use, including PMU placement, data quality, cybersecurity, and much more.

2.0 Brief Synchrophasor Definitions

This glossary summarizes key terms used in this White Paper. Mathematical details, subtleties and theoretical limitations of synchrophasor measurements are presented in Appendix A.

Synchrophasor	A synchrophasor is the result of a time-synchronized measurement of a phasor quantity. (The word synchrophasor is sometimes used, improperly, for the device that performs the measurement.)
Phasor	A phasor is an abstract, idealized representation of an electrical quantity such as voltage or current that is assumed to vary in time according to a perfect sinusoidal wave of constant frequency. A phasor contains two pieces of information: magnitude and phase angle. By convention, the magnitude is the root-mean-square (rms) value, or average height of the wave. The phase angle in degrees expresses a time shift of the sine wave relative to a reference clock. The difference between voltage phase angles at different locations is closely related to power flow across the network. This difference can be observed only if the measurements share the same time reference, thus “synchro-”. By definition, a phasor describes an entire cycle of a wave, not an instantaneous physical value.
Phasor measurement unit (PMU)	A PMU is a device that reports synchrophasors. PMUs also report frequency and ROCOF.
Frequency	The a.c. frequency describes how many complete reversals voltage and current undergo per second. Ideally, it is constant at the nominal value of 50.00 or 60.00 hertz (cycles per second). When frequency is not exactly constant, it is not trivial to define or measure.
Rate of change of frequency (ROCOF)	The rate of change of frequency, expressed in hertz per second, describes how rapidly the frequency is changing, indicating an imbalance between generation and load. In the ideal steady state, ROCOF would be zero. ROCOF is typically of interest during significant grid disturbances.
Total vector error (TVE)	TVE is a measure of the accuracy of a phasor that jointly accounts for errors in magnitude and angle. Existing PMU performance standards refer to a 1% TVE.
Sample/Report	A PMU takes many rapid physical measurements (samples) of voltage and/or current, computes phasor quantities from these samples, then time-stamps and reports the phasor for each cycle or two. The reporting rate is expressed in frames per second.

3.0 Distribution-Specific Applications and Use Cases

3.1 Overview

For synchrophasor applications of interest at the transmission level, algorithms typically compare measurements across large distances, even if the PMUs happen to be installed on distribution circuits—say, at substations, or plugged into 120-V wall outlets [1] [2]. Such analysis provides important insights for wide-area monitoring, including frequency and angle stability, grid oscillation modes and damping, or significant disturbance events [3] [4] [5] [6]. By contrast, distribution-specific applications will be primarily concerned with informing *local* decisions, usually comparing data from multiple locations behind the same distribution substation. For example, an algorithm may seek to determine the cause and effects of a fault on a distribution circuit based on voltage and current phasors along the feeder, while checking synchronized data from elsewhere to rule out a disturbance propagated from the transmission side.

The goal of synchrophasor-based applications is always to provide increased visibility and situational awareness, which can extend both above and below the substation. However, on the transmission side, PMUs augment extensive telemetry already in place for purposes of power system measurements (voltage and current magnitudes, real and reactive power) in near real-time. The higher resolution and precise time synchronization of synchrophasors crucially improves upon these conventional measurements, by revealing subtle changes and dynamics across the time dimension. These dynamics are seen in the relationships among quantities over large distances across the network, for example, in the context of oscillations and their damping. Across transmission systems, time-aligned frequency measurements, even without explicit phasor differences, can be highly informative [5], supplementing SCADA and EMS systems that approximate steady-state conditions.

By contrast, existing SCADA and customer meter data leave vast gaps of knowledge about distribution circuits, where even the physical properties and connectivity of the network itself are often in question. Many distribution system applications will need to first establish a baseline awareness of the distribution system's operating state, in order to detect and understand anomalies and problems. The goal will be to identify grid impacts arising from specific individual sources and how those affect the rest of the distribution feeder and the upstream system. This is difficult because the relevant signals to be measured are small, distribution systems are noisier than transmission systems in terms of data behavior, and algorithms must account for many variables. Moreover, these variables include unknowns such as exact impedances of the primary and secondary distribution lines, ABC phase connectivity, or operational status of rooftop photovoltaic systems.

3.2 Distribution Applications, Measurements and Data Needs

Many distribution-specific applications for synchrophasor (and more broadly, time-synchronized) data are discussed in [7], including:

- Event detection and classification
  - Voltage sage detection and analysis
  - Low- and high-impedance faults
  - Equipment health diagnostics
  - Fault location
- Topology detection
- Cyber-attack detection
- Model validation
  - Load models
  - Generator models
  - Phase (ABC) identification
  - Line segment impedances
  - Transformer and other device models
- DG characterization
  - Correlate feeder voltage changes with DG behavior
  - Detect reverse power flow
  - Disaggregate net metered DG from load
- Microgrid operation
  - Islanding
  - Load and generation balance
  - Resynchronization
- Distribution state estimation
- Phasor-based control

These applications vary substantially regarding their demands and requirements for PMU data. Specifically, they will have different needs for

- Temporal data resolution
- Absolute and relative measurement accuracy
- Communication volume, latency and continuity of data transfer
- Density and placement of PMUs.

A table of expected data requirements by application is reprinted from [7]. The quantitative table entries are not intended as conclusive statements, but reflect approximate engineering judgments that will likely be updated and refined as R&D advances and as field experience with these applications grows. The primary take-away is that different applications draw upon different physical principles and therefore have different measurement requirements.

Application	Measurement Quantities	Time Resolution	Accuracy	Latency & Continuity
Voltage magnitude profile & variability	Voltage magnitudes crucial, Voltage phase angle useful for recognition of tap changes	1 sec or better resolution is useful, synchronization between & among measurement locations critical	Changes in time of interest, absolute accuracy to 0.5% error adequate	Retain complete history
Awareness of real-time loads	Current magnitudes very useful, V phase angle can be proxy for current iff network impedances are known; current phase angle useful for P,Q decomposition & reverse power flow	1 cycle or better resolution reveals transient behaviors, full time domain characterization with up to 30 kHz sampling of interest to reveal harmonics	Absolute 0.5% error likely adequate	Operationally relevant latency on the order of 1 sec
Outage management	Voltage & current magnitudes	1 sec likely adequate	1% error likely adequate	1 sec latency likely adequate
System frequency & oscillation detection	Voltage phase angle essential	1 cycle or better & synchronization essential	Changes in time, not absolute accuracy of interest, 1% error adequate if stable	Retain complete history; latency requirement may vary, sub-second critical if informing protection
Island detection; Microgrid islanding & resynchronization	Voltage phase angle essential	1 cycle or better resolution	Insensitive to magnitude error, phase angle error stable to 0.01^{o}	Continuous monitoring, sub-second latency critical if informing protection
Distribution state estimation & SE-based topology detection	Voltage phasors; sensitive to placement & number of sensors; network model & load data important	Synchronization critical	Absolute accuracy on the order of 0.0001 p.u., requires correction for transducer errors	Operationally relevant latency on the order of 1 sec
Topology detection based on time-series signatures	Voltage phasors	1 cycle or better & synchronization critical	Changes in time, not absolute accuracy of interest, 0.5% error adequate if stable	Retain complete history, operationally relevant latency on the order of 1 sec
Topology detection based on source impedance	Voltage & current phasors	1 cycle or better & synchronization critical	Changes in time, not absolute accuracy of interest, 0.5% error adequate if stable	Operationally relevant latency on the order of 1 sec
(ABC) Phase identification	Voltage phase angles essential	1 sec or better for time-series approach; synchronization critical	Absolute accuracy of phase angle on the order of 1^{o} likely adequate	No particular need for latency or continuity
Model validation for line segment impedances	Voltage & current phasors	Synchronization critical	Absolute accuracy of all phasors is limiting factor, as good as 0.0001 p.u. for shorter segments	No particular need for latency or continuity
DG Characterization; Transformer, generator & load models	Voltage & current phasors	1 cycle or better reveals dynamic behaviors; synchronization between primary & secondary side of transformer critical	Changes in time, not absolute accuracy of interest, 0.5% error adequate if stable	No particular need for latency or continuity
Event detection & classification	Voltage & current magnitudes adequate for most events, phase angles useful	1 cycle or better, synchronization critical	Changes in time, not absolute accuracy of interest, 0.5% error adequate if stable	Continuous monitoring, operationally relevant latency on the order of 1 sec
Fault location	Voltage & current phasors	1 cycle or better, synchronization critical	Absolute accuracy of all phasors is limiting factor	Continuous monitoring, latency on the order of 1 sec
Phasor-based control	Voltage phasors	1 cycle or better	Absolute accuracy critical for steady-state optimization, but stable errors acceptable for disturbance rejection	Continuous monitoring, latency critical

Table 1: Expected Data Requirements for Various Types of Distribution PMU Applications [7].

Table 1 suggests that since some of the above tasks can be performed using measurements of lesser precision or resolution (and thus presumably lesser cost), it may be easier and more cost-effective to deploy specialized devices for specific tasks (using a common data network) rather than using a large fleet of standard devices to collect data for all distribution application tasks.[1]

For instance, a significant subset of applications to improve situational awareness of distribution systems (including event detection and real-time knowledge of loads and

[1] There is a ready parallel to transmission-level PMUs, which include M-class and P-class devices to collect data for executing different tasks.

generation) do not rely on the explicit measurement of the voltage phase angle shift between locations. In that sense, they don't need synchro*phasors* at all; they just need time-synchronized measurements of rms magnitudes, with synchronization on the order of a full cycle rather than a fraction of a degree. This means three orders of magnitude less precision in the time stamp (i.e. milliseconds instead of microseconds), which would greatly simplify the apparatus.

Also, numerous distribution utilities are investigating how much operationally relevant information can be leveraged from their advanced metering infrastructure (AMI) using smart meters of end-use customers. These approaches might involve some modification of firmware and data collection while working within the capabilities of the existing sensor and communication hardware.

These examples speak to the value of measurements other than high-precision synchrophasors. However, they still rely on a coordinated *network* of sensors. Also, they are consistent with the philosophy that a single network can serve multiple uses as well as multiple *user*s – for example, the billing, engineering, and operations departments within a utility. This suggests that the economy of system monitoring should be evaluated in a *collective* sense – that is, across the broadest set of supported applications – rather than in terms of individual, isolated use cases. The comparative cost of sophisticated sensor capabilities and crosscutting requirements such as installation, communications, data analytics, etc. remains a legitimate and important consideration. But there could be a false economy in designing a sensor and data infrastructure that precludes potentially valuable uses which are more demanding in data terms. As utilities know very well in the context of planning for load growth, a modest size increment in a capacity investment can have a large payoff if it avoids the need for additional investments a few years in the future. Likewise, it may be most cost-effective to build conservative 'headroom' into a sensing and data infrastructure at the outset.

On the other hand, a risk of over-designing a monitoring network is producing an excess of information: too many sensors sending too much data, driven by the time resolution and number of channels on each sensor. This could be counterproductive or disruptive for some uses. If finding and extracting the actionable intelligence from a voluminous or awkward data store requires any additional time or effort of the user, then more information may be harmful rather than helpful. This risk is more a function of the data management infrastructure than the sensor technology involved.[2] Any advanced grid data infrastructure should facilitate filtering and consuming the data at different resolutions and locations (i.e., near the sensor versus at some central location), as appropriate to the application purpose.

The following sections illustrate a subset of applications in more detail. They are not intended to provide exhaustive coverage, but convey a sense of some important emerging use cases.

[2] The Berkeley Tree Database (BTrDB) was developed to handle high-resolution time-series data from sensor networks specifically in light of this concern.

3.3 Event Detection, Classification and Cyber Attack Detection

High resolution measurements on a fast and readily searchable database provide many new opportunities for detecting events that previously could not be diagnosed, located, or observed at all. For example, Figure 1 shows an event observed with µPMUs on a 12-kV distribution circuit. SCADA did not register the transient current, and reported the increased steady-state current with a 7-second error in the time stamp.

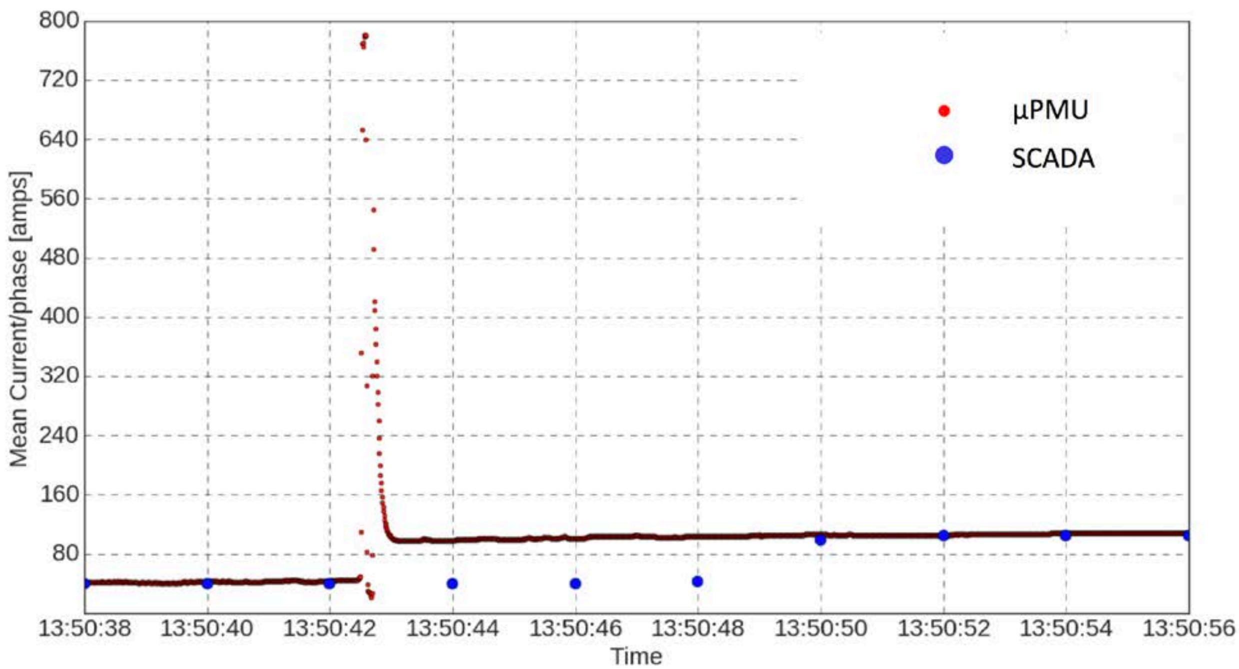

Figure 1: µPMU (120 frames per sec) versus SCADA (2-sec) recordings of a transient event. (LBNL)

Measurement synchronization is critical in order to identify a unique event and compare its magnitude as observed from different locations. As an example, Figures 2 and 3 illustrate an installation of three µPMUs on two distribution feeders, and a voltage disturbance event recorded by each µPMU. An event detection algorithm picked up the voltage sag at µPMU 3, where the magnitude dropped beyond a chosen threshold. Cross-referencing time stamped data from µPMUs 1 and 2 shows that while the event is recognizable at these locations, its per-unit magnitude there was much smaller and did not trigger an event report. From these observations, the operator or automated algorithm can deduce that the cause of the event is most likely located on Feeder 2 (below µPMU 3), since a more distant disturbance propagating through the transmission grid would have registered more equally at the two feeders.

Further analysis such as comparison of voltage and current measurements, or magnitude and angle data, can classify events as faults, abrupt changes in load or distributed generation, tap change operations, etc. The correlation, localization and classification of events with a set of data analytics based on µPMU measurements in the distribution grid is presented in [8].

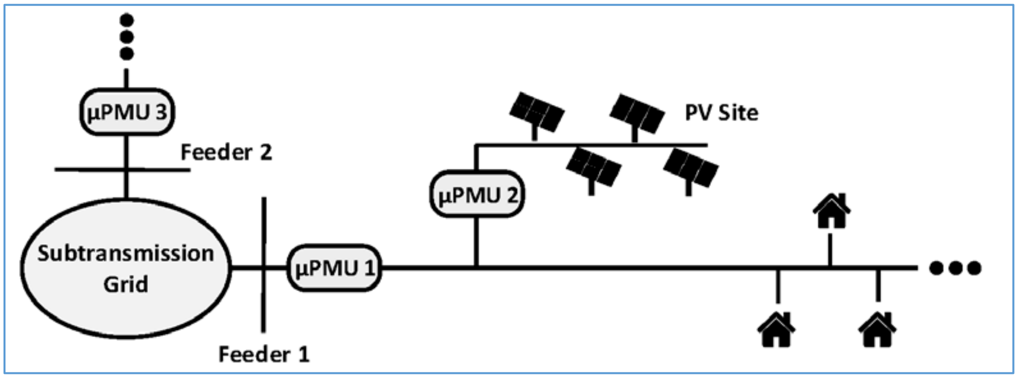

Figure 2: Schematic location of installed µPMUs on a 12-kV system, reprinted from [9].

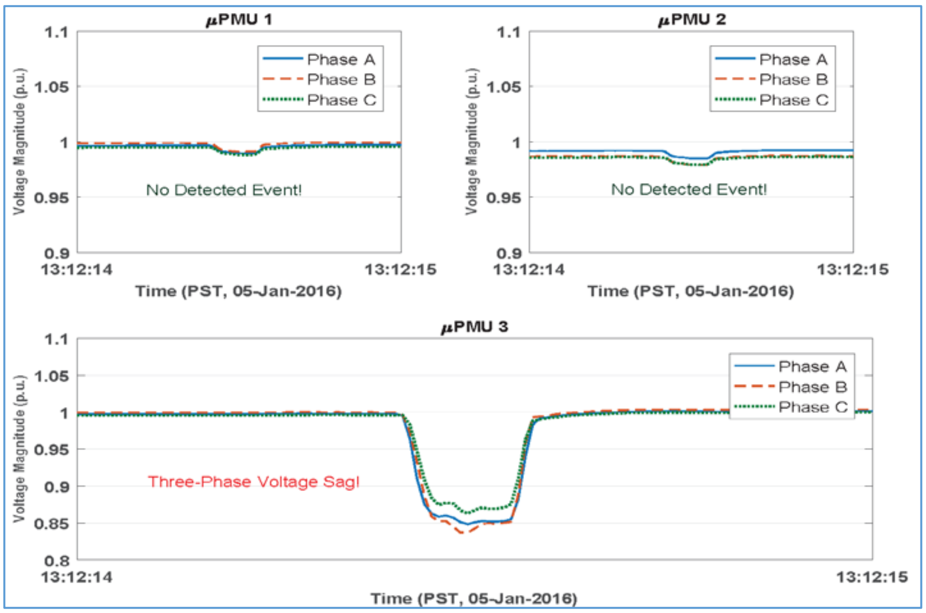

Figure 3: Voltage magnitude change in per-unit observed at the three µPMUs in previous figure. (LBNL)

Event detection and classification go hand in hand with cyber-attack detection. Cyber-attacks or attack rehearsals leave footprints in the physical measurements of the grid, though they may not cause obvious disruptions. For example, an attacker who hacked into a utility's SCADA system might test his ability to remotely actuate switches by opening and closing a switch that does not interrupt any load or trigger any alarm. This type of rehearsal would normally go unnoticed, but careful analysis of synchrophasor data could reveal the switch operation, and an automated comparison against known operating actions could flag the event. Conversely, detected anomalies can have cyber-related causes (such as unauthorized manipulation of equipment) or natural causes that may pose safety risks (such as high-impedance faults). In either case, early detection is essential to avert potential safety or reliability problems.

Early detection of cyber-intrusions is a crucial defensive strategy because even if an intrusion is not causing direct harm at present, it may be part of an attacker's learning process and pave the way for more serious future disruptions. Cyber-attacks can also defy N-1 security criteria, by compromising the redundancy in the grid that ordinarily limits the impact of any one element's failure. By the time any serious physical consequences such as loss of load are noticed, it may be too late to intervene.

To defend against cyber- and other threats with intelligence from the distribution system, an approach documented in [8] is to equip the grid with an event detector and classifier to first filter out anomalies, and then perform a forensic analysis of the likely source of the anomaly. This level of insight can be obtained by defining proper, searchable metrics as a function of synchrophasor readings. These metrics aim to capture specific characteristics that can be traced back to physical misuse of automation equipment, or data inconsistencies that indicate sensor malfunction or injection of falsified data.

Localizing and classifying events draws on knowledge about the system topology, physical laws and the implemented automation of the grid to fuse the data and test the hypotheses of a cyber-attack or other external factors as the probable cause. Rules are applied at different levels of data aggregation and may include inspection for anomalies in the voltage magnitude as well as fast changes in the current magnitude, active and reactive power and local frequency. Because a signature of disturbance events is to take the grid out of steady-state operation, the rules also check for validity of the steady-state equations using single or multiple µPMUs.

Figure 4 illustrates a hierarchical framework for analyzing grid data at different levels of aggregation. Next to each sensor, rules are applied to search for "anomalies" based on local readings only. Eventful data segments are annotated, stored, and prioritized for sending to higher levels of aggregation, where another set of rules correlate multiple sensor readings to identify a possible cyber-intrusion. An implementation of this architecture from [8] has each sensor "sandboxed" with a BeagleBoard; the database is hosted at http://powerdata.lbl.gov/explore.html. [10]

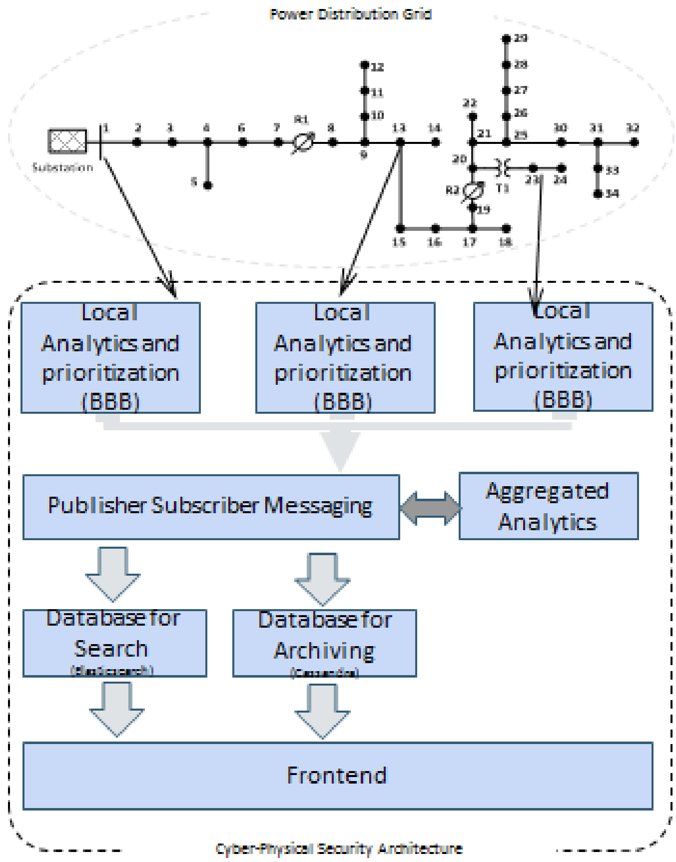

Figure 4: Implemented Hierarchical Intrusion Detection Architecture. (ASU)

3.4 Distribution Network Topology Detection

Topology detection algorithms help to determine the status of switches (open or closed) at known locations in the power system. Knowledge of the network topology is essential to maintain safe operations and estimate the system state accurately. The status of switches, and thus the network topology, is generally expected to be available through Supervisory Control And Data Acquisition (SCADA). In practice, however, switch status information may be sparse and outdated, because SCADA system poll instrumentation near these switches and the information is not synchronized. Although connectivity is not directly sensed by µPMUs, it can be inferred from phase angle differences between points on opposite sides of a switch.

Several approaches to topology detection based on µPMU measurements are described in the literature. For example, [11] is a model-less approach based on time-series data from a dynamic system. These data show specific patterns regarding state transitions, such as opening or closing of switches, as a kind of signature from each topology change. The proposed algorithm in [12] compares the actual time-varying pattern of voltage phasor values to a library of patterns associated with possible topologies of a given distribution network. The work in [13] proposed a method to reconstruct the interconnectedness of distribution networks based on dynamically related stochastic processes. The model-based approach in [14] uses a voting-based algorithm that looks for the minimal difference between measured and calculated voltage angle or voltage magnitude to indicate the actual topology. Theoretically it can be shown that accurate topology estimation for radial distribution grids is possible with measurements of bus voltages limited to a sparse set. This is due to specific trends that the radial topology imposes on second-order voltage statistics, which enable the relative locations of distribution grid buses to be determined [15]. In sum, the problem of network topology detection with high-precision synchrophasor data is an active area of academic investigation, which has both intrinsic mathematical appeal and practical significance.

A simpler but highly relevant case of topology detection is to recognize when a conductor is severed (for example, due to a falling tree limb). Just like an open switch, a broken conductor will instantly cause a pronounced phase angle difference across the break point that is easy to detect with PMUs. This information has proven capable of actuating protection systems before an energized conductor would hit the ground and present a serious injury or fire hazard [16].

Another related use case is the detection of changes in the connectivity status of distributed generation (DG) through analysis of data from only the utility side of the interconnection, if direct telemetry is unavailable. Online monitoring of DG for real-time detection of unexpected trips or re-connection may identify protection coordination issues on the circuit. Knowing how much generation is actually connected at any given time will add confidence to other operations. This application is also related to DG-Load disaggregation, discussed below.

3.5 Model Validation and Phase Identification

Correct distribution circuit models are most important in the planning context, for example, when predicting the impacts of variable distributed resources. A comprehensive model would incorporate information about the impedance and connectivity of each electrical component in the network, including details about the physical layout and geometry (e.g. spacing of underground cables), time varying behaviors (e.g. load profiles and dependence on voltage), and response to disturbances (e.g. dynamic inverter models). In practice, it is difficult to compile and verify the accuracy of all this information, and planners must work with estimates that may or may not reflect how the actual system will behave.

Precise voltage and current phasor measurements can be used to compute impedances of line segments and other components within the three-phase network, and compare empirical values to model data [17]. This application is very sensitive to data quality (specifically, transducer errors) and might benefit from further development (e.g. using suitable mathematical regression methods [7]). However, short of validating a circuit model in complete detail, one straightforward and immediately useful aspect is to identify the phase connectivity.

The objective of phase identification is to recognize, track and report the connectivity and loading of the three phases (A,B,C) throughout the distribution system so as to prevent excessive imbalance. Distribution system maps and models tend to lack reliable information about the phase connection of single-phase laterals, double-phase laterals, or individual customers on the three-phase main feeder. This connectivity information is not entirely static, since during restoration work such as repairs after a major storm, the phasing may be changed, either deliberately or by mistake. Historically, utilities have relied on manual notations about connectivity from field crews performing the work. Smart meters are now a means for checking the loading of main and laterals, but determining the load by phase still depends on correct knowledge of the topology.

Unlike transmission systems, which benefit from the statistical aggregation of a large number of customers connected to different phases, distribution system loads can be very unequal. Load current imbalances on the order of 10% are not uncommon.

Because of finite source impedance, unbalanced currents will furthermore result in unbalanced voltages on the distribution network.[3] This can physically damage three-phase motors, and may also interfere with the controls of three-phase inverters. ANSI standard C84.1 specifies that balanced voltages be within 3%. NEMA motor ratings require a more stringent 1% voltage unbalance [18].

Increasing penetration of solar PV generation and electric vehicles brings an increased risk of distribution system unbalance. For example, a small number of customers with EVs on the same phase can suddenly have a greater impact on unbalance. At the same time, equipment based on inverter technology can be more sensitive to unbalanced voltages.

Power system protection equipment (i.e. relays, circuit breakers, reclosers) and voltage regulators also may not operate as designed under significant phase imbalance. Their misoperation can cause nuisance tripping and voltage magnitude violations, respectively. Finally, appropriate phase balancing (i.e. within 3%) is necessary to maximize asset utilization (e.g. balance transformer capacity across all three phases), minimize energy losses, and help prevent equipment issues in the distribution network.

Conventional methods for phase identification either rely on data from Geographic Information Systems (GIS), which is a known source of errors since it requires mapping by personnel, or commercially available dedicated phase identification tools. These tools require the installation of phase reference nodes, and the method is also impacted by transformer (delta-wye) phasing. Some tools rely on verification rules or clustering based approaches on time-series of voltage or consumption readings [19].

By contrast, PMUs measure voltage phase angles directly, which offers immediate visibility of phases. These phase angles will be 120^o apart on the three phases, with a very small separation (on the order of a degree) between points on the same phase along a distribution circuit. Each delta-wye transformer or lateral tap introduces a 30^o phase shift between locations, which makes the phase association much less obvious. In this situation, correlation between the time series signatures of voltage magnitude and/or angle can be used as an additional identifying tool. This is easiest during a large asymmetrical disturbance, like the one illustrated in Figure 5.

[3] Unbalanced currents and voltages are also known as negative-sequence or zero-sequence components, which add to the positive-sequence quantities that represent balanced three-phase operation.

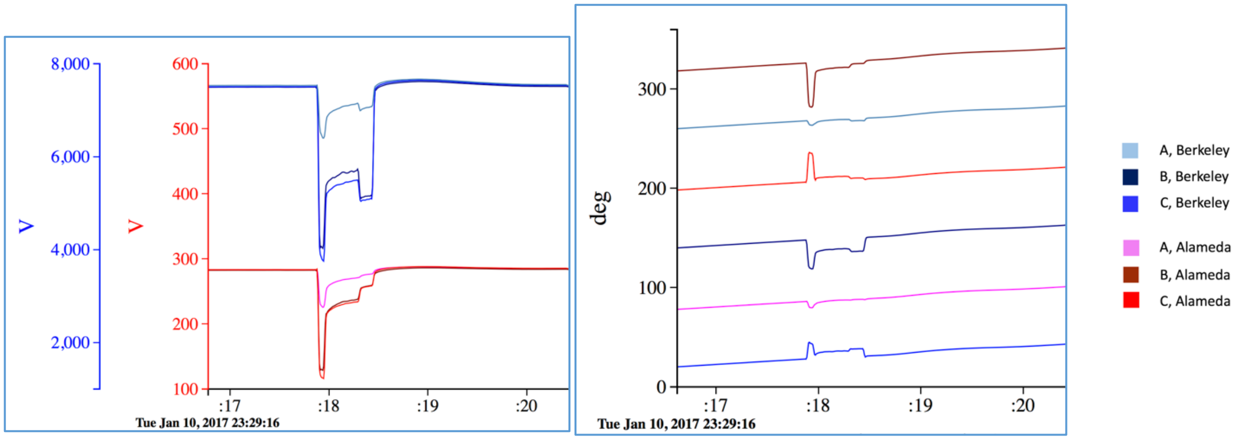

Figure 5: Voltage magnitudes and angles during a disturbance event at two locations separated by the transmission grid and multiple delta-wye transformers. Graphs show µPMU data on BTrDB, with time axis in seconds. (UC Berkeley, LBNL, Power Standards Lab)

The example in Figure 5 shows voltage magnitudes and phase angles (relative to the same clock signal) during an event observed at two locations, Berkeley and Alameda. These two Bay Area locations are separated by a 115-kV transmission network and multiple transformers. Without any network model information, the phasing can be identified with confidence based on the different shape of the angle disturbance for each phase (right). This is possible despite the fact that matching phases are shifted by 180◦. It is also possible to match Phase A between the two locations based on the smaller per-unit magnitude of the voltage sag (left), but the association between Phases B and C on the basis of voltage magnitude alone is much less conclusive than inspection of the angle measurements.

Using PMU data obviates the need for specific equipment to actively inject a signal for phase identification. Phase identification relies on the comparison of time-synchronized voltage phase angle measurements, but is not sensitive to their absolute accuracy.

3.6 Equipment Health Diagnostics

High-precision synchrophasor measurements can detect early signs of equipment aging, mis-operation or impending failure from the electrical signature. This detection can help prevent costly damage or outages, by taking the equipment out of service for repair or replacement before it fails.

Aging and deterioration in equipment such as distribution transformers or switchgear can be difficult to diagnose inexpensively while the equipment is online. Smaller service transformers (e.g. 25 kVA) that supply only a few customers are typically just replaced upon failure. Larger capacity equipment such as substation transformers can be tested using dissolved gas analysis (DGA) of the transformer oil that reveals chemical evidence of degradation (specifically, of the insulation material) and provides some estimate of remaining transformer life.

In a pilot deployment of µPMUs at Riverside Public Utilities (funded by ARPA-E), researchers from Lawrence Berkeley National Laboratory discovered a voltage sag that characteristically followed within several cycles of tap change operations, accompanied by a small current transient. The signature is shown in Figure 6. This observation prompted utility personnel to perform a field inspection of the transformer, which revealed an oil leak that could have resulted in costly damage [20].

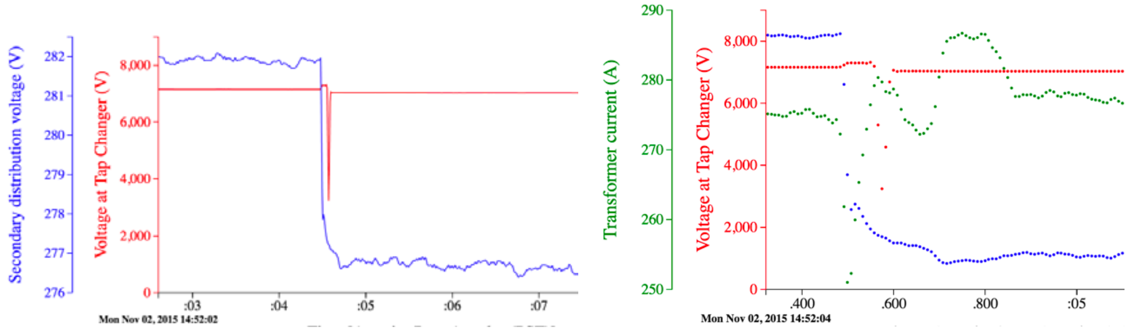

Figure 6: Anomaly in the tap change signature gave early warning of deterioration at a substation transformer. PSL µPMU data visualized in Berkeley Tree Database plotter with horizontal axis in seconds; right graph shows individual 120-Hz data points. (UC Berkeley and LBNL)

Ongoing condition monitoring of utility equipment would help both to prevent specific device failures and to establish a general improved knowledge base for planning purposes.

Equipment health diagnostics benefit from precision time-series measurements, time granularity of data on the order of a cycle or better, and synchronization of measurements made across different locations for validation through cross-referencing. There is no specific threshold for absolute accuracy. Explicit voltage phase angle data may be useful but are not intrinsically necessary.

3.7 DG-Load Disaggregation

Disaggregation of net metered distributed generation (DG) from customer load separates the load demand from the amount of generation provided by the DG at any instant. It uses high-precision measurements on the utility side of the meter to estimate the actual generation and amount of load offset behind the meter.

When the distribution utility lacks access to separate load and generation telemetry from customer premises, net metered solar generation obscures or “masks” an unknown amount of load from the system operator’s view. This “masked load” implies a greater system exposure to contingencies. The masked load must be accounted for to assure adequate generation reserves in case of simultaneous tripping of many DG units, and to assure safe cold load pickup during system restoration following an outage.

Innovative algorithms combine PMU measurements on the utility side with available solar irradiance data for a high-fidelity estimate of individual PV generation and masked load, even when not directly metered.

The conventional approach to estimating DG output is based on reported generation capacity, irradiance data, and generalized solar production models. These models lack specific information about operational up-time or degradation of any given PV installation, and the model fidelity suffers during variable (non-clear-sky) conditions, which are precisely those of greatest interest. Alternative approaches to identifying actual DG output and masked load include adding telemetry on customer premises, or accessing customer PV generation data by way of third-party online platforms; each of these have intrinsic technical and non-technical challenges.

In addition to revealing masked load, visibility of short-term DG behavior from the utility side of the meter offers the potential benefits of validating smart inverter performance, and correlating PV generation with feeder conditions such as voltage profile changes.

DG-load disaggregation draws on continuous high-resolution time-series measurements of voltage and current, including local current vs. voltage phase angle for displacement power factor, from a single measurement point. There is no critical threshold for sampling rates or absolute accuracy; the fidelity of the estimate will improve in direct relation to data quality. Extant algorithms successfully used µPMU data with 120-Hz sampling and accuracy constrained by revenue-grade instrument transformers. This use case does not require comparative voltage phase angle measurements between locations. Future analytics might be enhanced by algorithms that use simultaneous power quality measurements in the time domain to identify harmonic signatures.

Such signatures, if developed from aggregated injection/load measurements samples at high rates, can lead to individual load identification above and beyond load-DG disaggregation. Simplified tests based on ramp rates or change in power factors can identify equipment usage (say, air conditioners or dishwashers) within an aggregate energy profile for a household obtained from a smart meter [21] [22] [23] [24]. Use of µPMUs with far greater sampling rate can improve over the performance based on 15-min AMI data used in past work.

Development of non-intrusive load identification schemes, however, has led to important questions regarding user privacy and security. In particular, it is known that smart meter time-series can be used for occupancy detection for households [25]. For high-fidelity µPMUs with richer data, these concerns warrant further analysis. In particular, introducing some randomness (either in measurement or in consumption profiles through the use of small storage) has been suggested [26] to avoid exact disaggregation and its associated privacy issues.

3.8 Fault Location

The vast majority of faults in power systems occur at the distribution level. The traditional method of locating a fault on a distribution system is for a utility employee or crew to travel along the feeder where a protective device has operated, or where customers have reported an outage. The crew searches for the exact fault location visually in the case of overhead lines (e.g. by looking for an operated cutout fuse), or with an underground fault locator.

Synchrophasor measurements can complement physical inspection to reduce both outage duration and cost to the utility. A simple approach is to estimate the fault location using measurements of the fault current and voltage at the substation and a feeder impedance model. However, this method gives no unique solution, since multiple combinations of fault location and impedance may yield the same measurements at the substation. Also, the contribution of distributed generation (DG) to fault current, which is not directly observable from the substation, can lead to errors in fault location.

The accuracy of fault location can be improved by collecting synchronized measurements of current and/or voltage at additional locations throughout the distribution system. Several approaches based on search, optimization, and state estimation techniques have been demonstrated to locate faults using distributed measurements.

Many faults have been detected by µPMUs on distribution networks. An example is shown in Figure 7 [27]. As a first approximation, comparing the currents and voltage drops measured by multiple µPMUs distributed over a geographic area can indicate which of the µPMUs are upstream versus downstream of the fault location. It is also evident from the three-phase µPMU data which phases are affected by the fault. The specific fault location can then be computed from the set of measurements using a network model. Some algorithms, using simulated synchrophasor data, have demonstrated techniques for fault location to within a few meters or a few tens of meters [28] [29] [30]. Although these algorithms have yet to be validated using real µPMU measurements, the accuracy remains promising when simulated measurement errors are added [29] [30].

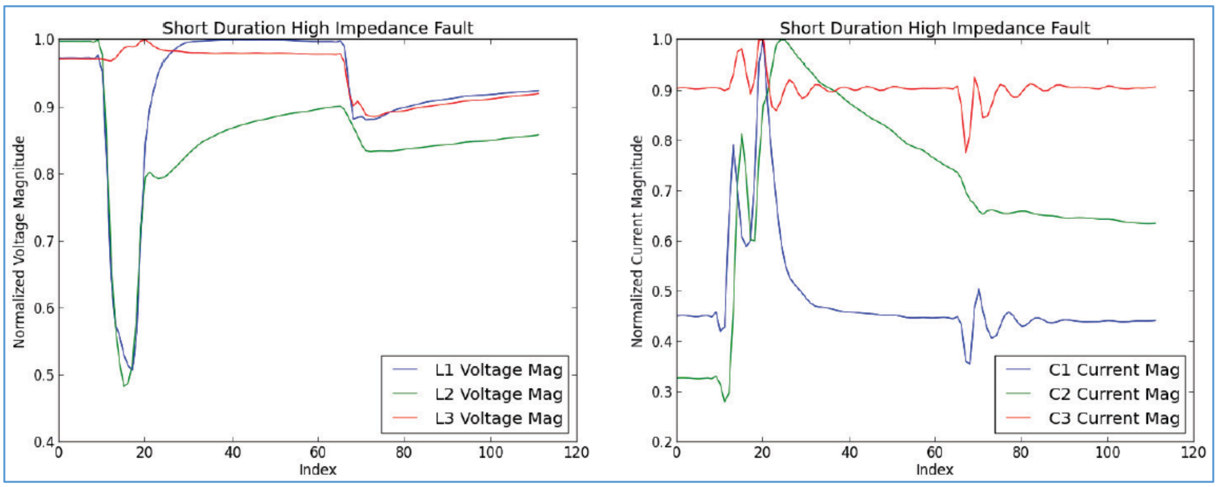

Figure 7: Three-phase voltage and current magnitude data streams collected by a µPMU during a high-impedance fault [27]. The time window displayed is one minute.

Fault location requires synchronized time-series measurements with time resolution at least on the order of a cycle (1/60 sec). While fault location methods using only voltage phasors [29] and only current phasors [30] have been demonstrated, maximum location accuracy ideally uses both magnitude and phase of both voltage and current. Since distributed generation can contribute significantly to fault current, measurements from the points of common coupling of DG installations are particularly advantageous. Knowledge of the errors introduced by current and potential transformers under high current and low voltage conditions is helpful to determine the uncertainty and confidence level in fault location.

3.9 Network-Level Steady State Estimation

State estimation generally means reconciling available physical measurements (of imperfect accuracy) with mathematical relationships (based on an imperfect model) to obtain a best-fit estimate of the state variables (voltage magnitude and phase) at each network node. A rigorous state estimation requires at least as many measurements as there are nodes in the network.

This condition is readily met in transmission systems, where a typical node represents a substation, but is much more difficult to achieve in distribution systems, where every service transformer is a node (i.e., branch point in the circuit). Without instrumentation at every node, some extrapolation from available data is possible, for example, by treating historical information about loads as "pseudo-measurements" [31]. The success of any state estimation methods is sensitive to the absolute accuracy of phasor measurements, including transducer errors.

Short of providing complete knowledge of a network's exact operating state at every moment in time, however, estimation of the steady state has been shown useful to identify impending unstable regimes like voltage collapse and bifurcations [32] [33] [34]. Time-synchronized, high-fidelity measurements from distribution PMUs can identify dynamic deviations from normal stable conditions. PMU data can also be used to estimate the network steady state using regression over known covariance statistics. Time-stamped measurements are particularly helpful in this regard as they enable measurement of `delayed' covariances at observed buses [35]. This significantly improves the estimation of steady state behavior and ambient noise statistics over traditional methods.

4.0 Synchrophasor networks and implementation

A number of technical considerations affect the deployment and use of distribution-level PMUs and other time-synchronized monitoring systems. This section addresses those, starting with a discussion of distribution-level monitoring devices and the key issues around data quality and availability.

4.1 Distribution-level PMUs and monitors

PMUs designed specifically for use in distribution systems are sometimes referred to as D-PMUs, or as micro-PMUs (µPMUs). The term µPMU originated from an ARPA-E funded project [36], but it is not trademarked and was conceived as a generic term to describe PMUs with extremely high measurement precision.

The ARPA-E µPMU achieves a measurement resolution of 0.0001 per-unit voltage magnitude and 0.01 degrees of angle, which exceeds by several orders of magnitude the typical resolution of PMUs used in the transmission context [33]. The motivation for developing such a device is that voltage and current measurements in distribution systems are generally characterized by a smaller signal-to-noise ratio than in transmission systems. In particular, the change in voltage phase angle along a distribution circuit with typical power flows and impedances is on the order of fractions of a degree, compared to full degrees or tens of degrees along a transmission line [7].

As noted above, not every application for synchronized distribution system measurements draws on the explicit phase angle difference at different points on the circuit, so it is not strictly necessary for every sensor to be a distribution-specific PMU. On the other hand, there is likely an economy of installing standardized devices in a routine manner that by default are capable of supporting any and all applications, including the most exacting ones.

PMUs for use in either transmission or distribution settings are commercially available from various manufacturers, and prototypes have been built for research purposes at several universities. Their characteristics vary and are they are too diverse to review here.

A key distinction lies between "distribution PMUs" that are installed on the distribution system for purposes of diagnosing phenomena on local distribution feeders, versus those effectively looking up into the transmission system. In the first case, which is of primary interest to this Task Team and this report, distribution synchrophasors are compared between and among PMUs below the same substation. In the second case, PMUs are connected at the primary or secondary distribution voltage for convenience of access, but their synchrophasors are interpreted alongside data from more distant locations. Distribution PMUs may make a useful contribution for transmission level uses in this second way, but the specific requirements for precision and data quality from distribution PMU apply to the first category.

PMU capabilities are already embedded in many devices such as relays and digital fault recorders, used in the transmission context and installed at some distribution substations. An important question for distribution PMUs is whether they can be embedded in common devices such as transformers, inverters or protective devices at reasonable cost, so that their deployment will not require a specific labor effort.

4.2 Data quality and noise

High-resolution and high-precision synchronized measurements can support sophisticated data-driven methods for improved management and operation of power distribution systems. However, the performance of a PMU sensor itself does not guarantee data quality. Impairments can be imposed on synchrophasor data during acquisition (at the level of the PMU or the instrument transformer), communications (at the network layer), and processing (e.g. phasor concentration). Noise, latency, data loss and timing issues are notable impairments that can affect the quality of synchrophasor data and reduce the performance of data-driven applications including control, monitoring and protection systems.

Power distribution networks generally exhibit more pronounced and erratic variations of operating parameters than transmission systems. At the transmission level, statistical aggregation tends to smooth out time-varying loads, weather, and phase imbalance. In addition, noise and higher frequency components of high-resolution measurements at distribution networks might come from various dynamic behaviors of inverter and power electronic based components including DG, EVs, and smart and energy efficient loads. There are not yet extensive studies or practical guidelines to distinguish, classify and label between signal and noise at different time scales in distribution system measurements.

In the context of calculations with synchrophasors, it is important to distinguish between two types of noise: namely, *input noise* and *data noise*.[37] In real-world scenarios, the input voltage or current signals that enter the input terminals of a PMU may be corrupted by additive or multiplicative measurement noise. Instrument transformers are an important source that is discussed in the following section. Moreover, harmonics and sub-harmonics can exist in the spectrum of these input signals, especially near inverter-based sources. Electromagnetic interference, caused by lightning and nearby wireless devices, can also aggravate signal distortion. The superposition of the above distortions and noises in the input voltage or current signals is called *input noise*. Analogous to stochastic systems, the input noise manifests itself at the output of the PMU as *data noise*. Data noise can be thought of as the random uncertainty in the reported synchrophasor data.

The level of the input noise depends on several factors, including the type of environment in which the PMU is installed (for example, inside a substation building or mounted on a distribution pole). In practice, phasor measurements in distribution grids are affected by various errors and white noises to a greater extent than those measurements in transmission systems. Some distribution network sensor locations are harsh in terms of electrical noise, where the quality of PMU data can be significantly compromised. Most methods that aim to detect bad PMU data rely on data history to identify low-quality measurements, and will therefore be imperfect. A high level of input noise may overwhelm the signal and render the reported synchrophasor data useless for certain applications.

When judging data quality, context is of the essence: what, exactly, does the user wish to infer from the data, and do with that information? It is worth noting that input noise may contain valid information that simply does not fit the context in which it is observed and interpreted. For example, harmonic content would be considered noise when calculating the phasor of a presumed pure sinusoidal signal at the fundamental frequency, but the same harmonics might represent the signal of interest in a time-domain power quality measurement.

For the purpose of data analysis, the difference between the reported data and the actual state of the grid can generally be modeled by data noise. As illustrated in Figure 8, the statistics of the data noise depend on how the input noise propagates through the system (in this case, the stochastic system is a PMU). Specifically, the level of the data noise depends the level of the input noise, the low-pass filtering process, and the phasor calculation method.

Since different PMU vendors employ different filtering and phasor calculation methods, there is no unique solution for the explicit analysis of data noise. This analysis might be facilitated in the future either by full transparency of the various filtering and calculation methods, or by means of big data analysis drawing on many PMUs of different types.

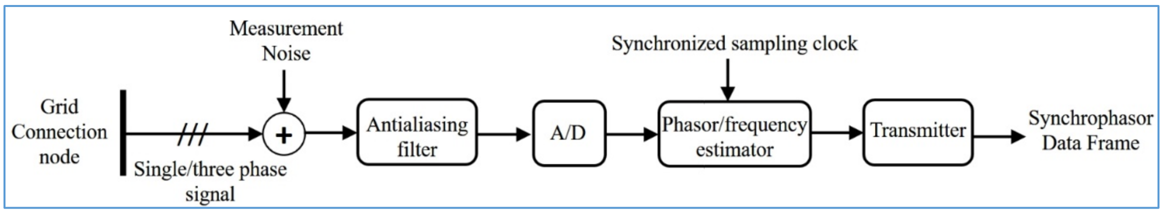

Figure 8: The generic procedure of synchrophasor data acquisition with additive input noise. (ASU)

In sum, synchrophasor data end users deal with data noise, not the input noise. This assumes that communication systems do not impose any uncertainty on the synchrophasor data, i.e., the received synchrophasor data frames at the destination are identical to the transmitted data frames.

Field data from actual PMUs deployed in the grid are subject to more uncertainties than simulation data created by software models. It is possible to make simulated data more realistic by adding a deliberate noise component (for example, before testing an algorithm on this data). Here, a crucial question is whether the data noise in the actual field data follows a Gaussian (i.e., normal) distribution.

To find the answer to the above question, extensive analysis of field datasets is very helpful. The noise characteristics of the electrical signals on the distribution level have been studied in [38] and [39], by collecting instantaneous a.c. data and analyzing the frequency spectrum. This work found that the signals overall have about 60 dB to 70 dB noise, and there exists more noise in the low frequency band. Monitoring and control applications should adopt different processing methods depending on the distribution of the data noise. For example, a Kalman filtering method is the optimal choice for dynamic state estimation when the state and measurement are jointly Gaussian random processes, whereas a particle filtering method is a better choice for dynamic state estimation under non-Gaussian data noise assumption [40].

4.3 Application Sensitivity to Data Quality

As noted earlier, different applications exhibit very different sensitivity levels with respect to data noise. Some of the most sensitive applications are summarized below.

- State estimation and networked control applications. State estimation likely has the highest noise sensitivity. If operating decisions are made on the basis of inaccurate estimates of system's parameters, they may have unintended consequences and can lead to system instability. Accurate estimation of static and dynamic power system states is also vital for the performance of networked control systems. A dynamic state estimator should be highly robust against data noise, and PMUs with the greatest noise rejection capability should be used. However, special attention must be paid to noise-delay tradeoff in data acquisition, since a higher noise rejection is translated into a greater delay, which is detrimental to control functionality [34].
- Network topology and disturbance detection applications. Monitoring of voltage and frequency disturbances is crucial for early detection of system instability, while monitoring of transients after changes in switch status (open or closed) helps detect network topology. Data noise may affect the decision-making process and result in errors, false alarms, and potentially unwanted control actions. Most monitoring applications will likely require using M-class PMUs, which deliver high data accuracy at the expense of longer calculation delays.
- High-level control applications. Hierarchical and droop-based control applications in active distribution networks and microgrids can benefit from synchrophasor datasets. In comparison with networked control applications that rely heavily on real-time communications, high-level applications are less sensitive to data noise. However, data noise may undermine the quality of control process through misleading information extracted from synchrophasor datasets. The underlying premise is that high-level control applications are usually event-triggered and operate over longer time span. Hence, the network has enough time to diminish the impact of noise by post-processing of datasets.
- Fault detection applications. Protection applications are highly time-critical. In situations where protective devices will operate based on synchrophasor information, the noise rejection capability of PMUs can be sacrificed in favor of delay minimization, as in P-class instruments. However, data noise may result in inaccurate estimation of the fault parameters such as fault impedance and location.
- Line Parameter and Thévenin Source Impedance Estimation. Line impedance estimation relies on small differences between measurements and is therefore highly sensitive to errors. The voltage drop over a distribution line segment (i.e., the signal of interest) may well be smaller than the uncompensated transducer errors, leading to an unacceptable signal-to-noise ratio. Likewise, the Thévenin source impedance estimated from PMUs at a substation can have a high error variance across different windows of data samples, and thus high uncertainty. Because these are off-line calculations, time is not of the essence, and any available processing techniques for improving data quality can be brought to bear.

4.4 Transducer Errors

PMUs generally receive their input through couplers or transducers, namely, current transformers (CTs) and potential or voltage transformers (PTs or VTs) that bring grid current and voltage levels to a tolerable range for the PMU device.[4]

These transducers become a significant source of both magnitude and angle error when interpreting the synchrophasor output. Specifically, the relationship between voltage or current magnitudes on the primary and secondary side of the instrument transformer will depart slightly from the ratio reported on the name plate, for reasons that include aging, environmental effects, and the burden on the instrument [41]. Signals also get phase-shifted when passing through the complex impedance. Individual PTs and CTs vary, even if they are of the same make and model, and even among the different phases (A,B,C) at the same location.

Transducer errors are typically reported by a complex value called "ratio error," a multiplicative term that corrupts the reported phasors. IEEE Standard 57.13 specifies the level of allowable error that a transducer (PT or CT) can introduce into a current or voltage measurement taken from its secondary side. The standard defines a set of transducer accuracy classes, each of which has an "error parallelogram" that specifies bounds on simultaneous magnitude and angle error.[5] IEEE 57.13 has been in place since 1968, and the parallelogram approach to defining transformer error is well-established.

As discussed further in Appendix A.2 on total vector error, there is some intuitive difficulty in reading an error parallelogram because of the different nature and representation of magnitude and angle quantities. Magnitude error, on the vertical axes of Figure 9, is given in terms of a ratio correction factor (RCF). The RCF is a ratio of transformer turns ratios, defined for an individual transformer as the ratio between the true turns ratio of that transformer and the transformer model's nominal turns ratio. In other words, the primary-side current or voltage of a particular transformer is equal to its secondary-side value multiplied by its model's nominal turns ratio, multiplied by that transformer's individual RCF.

Angle error, on the horizontal axes of Figure 9, is given in terms of "transformer phase angle" (generally represented as β for CTs and γ for PTs). β or γ is defined as the number of degrees or degree-minutes by which the transformer's secondary current or voltage angle leads its primary.

---

[4] For example, a µPMU can safely receive an input up to 690V, which obviously needs to be stepped down from measuring, say, a 12-kV bus. Typical CT outputs for grid installations are on a 5-amp scale. PSL's µPMUs use an additional ultra-high-precision CT to further transform the 5-A current signal into a voltage signal on a 0.33-V scale; this adds negligible error compared to the original high-power CT.

[5] The rationale for the parallelogram can be understood within the traditional context, apart from synchrophasors, where the principal information to be gleaned from high-precision instrument transformers is the real power P transferred at the local metering point. Since P depends on the cosine of the angle difference between voltage and current, and since current angle usually lags behind voltage, either a low reading of voltage angle or a high reading of current angle would increase the computed value of P. Thus, the impact on power measurement of a low voltage magnitude reading (requiring an RCF > 1) tends to be compensated by a lagging error in the voltage phase angle, whereas a low current magnitude reading is compensated by a leading error in the current phase angle. Similarly, high voltage or current magnitude readings would be compensated by the opposite angle errors.

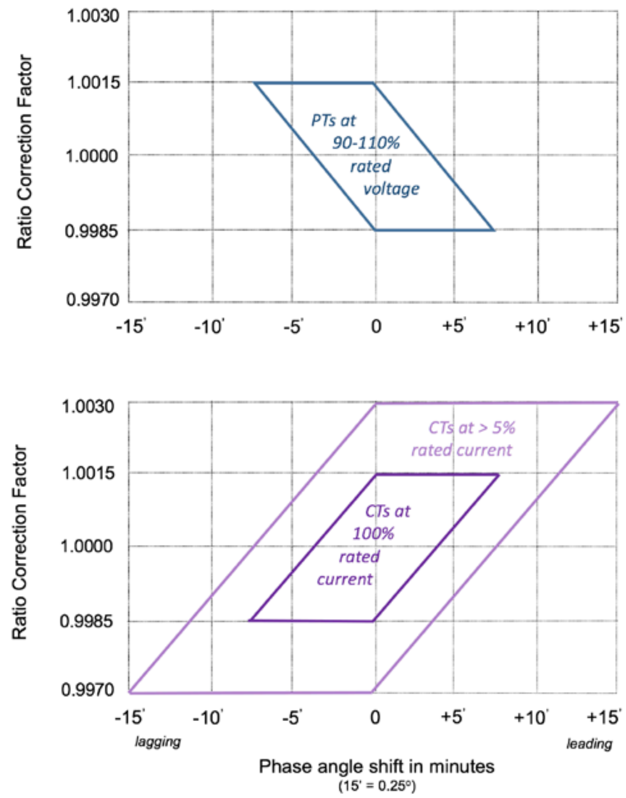

Figure 9: Error parallelograms for Class 0.15 instrument transformers, showing permissible error bounds for potential transformers (upper) and current transformers (lower graph). (Adapted from IEEE 57.13.6)

Instrument transformer classes are named for their RCF values. For a transformer of Class 0.15, with parallelogram pictured in Figure 9, the RCF must lie between 0.9985 and 1.0015. As can be seen in the figure, its phase angle error is bounded at approximately 7.5 degree-minutes.

To date, distribution-level PMUs have mostly been deployed on the secondary sides of Class 0.3 CTs and PTs, with RCF values constrained to lie between 0.997 and 1.003 and β or γ values bounded at approximately 15 minutes.

Note that the description and diagrams above apply only to metering specifications. Another class of specifications, which covers the use of transducers in relaying, defines acceptable transformer behavior during overcurrent and abnormal conditions.[6] These relaying specifications have a different nomenclature from the metering classes discussed here.

There is also an important difference between specifications for PTs and CTs. PTs are required to stay within the bounds of their error parallelograms only so long as their primary side voltage lies within 90% to 110% of the PT's rated value. CT's, on the other hand, are required to stay within the bounds of their error parallelograms even if primary currents drop as low as 10% of the rated value. Below 10% of rated primary current, the error constraints are relaxed. Between

[6] While metering applications are concerned with precision and accuracy during normal operating conditions, transducers for protection systems must remain dependable under abnormal conditions such as large fault currents, even if the measurement values are approximate. A similar distinction applies to measurement (M-class) and protection (P-class) PMUs.

5% and 10% of rated primary current, the allowable values of magnitude and angle error double, effectively doubling the size of the error parallelogram, and below 5% of the rated primary current, there are no longer any limits on the error that the CT is allowed to induce per IEEE 57.13.

These requirements make sense in that the operating voltage of a distribution circuit is expected to vary much less than the load. During normal operation of a distribution circuit, keeping voltage within 10% of a PT's rated value should not be a concern. However, CTs could certainly experience primary currents of less than 5% of rated values, making them a potential source of large error during low-load periods.

An important area of investigation is the drift of transducer errors over time. Stable transducer errors would not impact any application algorithms based on comparing measurements at different times. However, drift in transducer errors would introduce an apparent change in the measurements. Future research may examine what explanatory variables (such as transformer burden or temperature) affect transducer error, and whether these causes produce stable, predictable or unpredictable bias in PMU measurements.

Initial observations from distribution synchrophasors suggest that medium-voltage PTs and CTs in the field may operate near the edge or even outside the error parallelogram, and that caution is therefore indicated when interpreting high-precision PMU measurements that use these transducers for their input [17]. It has been reported in [41] that the ratio errors can be as large as ±6% for voltage magnitude, $\pm 4^{o}$ for voltage angle, ±10% for current magnitude, and $\pm 6.67^{o}$ for current angle. Such a level of error would deteriorate the results of many downstream applications.

It is worth noting that before the advent of synchrophasors, transducer phase angle shifts in either voltage or current measurements would have had little direct meaning or consequence *per se*, being significant only by way of distorting overall voltage, current and power quantities. High-precision synchrophasor data analysis may herald a new level of scrutiny of instrument transformers and their performance in both theory and practice.

4.5 PMU Placement: Theory

The placement of PMUs is informed by two very different sets of considerations: theoretical and practical. The theory is concerned with how much information about the electrical network is conveyed by PMU data from particular locations, and how to maximize this information, assuming that PMUs are scarce.[7] PMU placement strategies studied in the literature primarily aim to guarantee that the *state* of the grid (uniquely expressible as the voltage magnitu3de and phase angle at every node) can be estimated with a minimum number of PMUs.

Some techniques for identifying optimal PMU placement ensure observability (in the rigorous mathematical sense) by ensuring the algebraic invertibility of linearized load-flow models (e.g., [42] [43]), while some others focus on the topological observability (e.g., [44] [45]). Another approach to the placement problem is to impose constraints that avoid the formation of *"critical measurements*," or bottlenecks in the mathematical state estimation process (e.g., [46]).

While some methods rely solely on the PMU measurements to guarantee observability [47], others have proposed the use of hybrid state estimation, in which the placement of the PMUs and SCADA meters is considered within the same framework, and input data from both sources are used to obtain the state of the grid. If the cost of PMUs or new sensor installation is high, this approach could be an economic way of utilizing PMUs while benefiting from SCADA data that is already available [48].

The objective of full observability is much more difficult to meet in the distribution than the transmission context, since there are many more nodes or branch points on a distribution circuit (essentially, each service transformer is a node), and a complete state estimation requires at least one measurement per node.

But not all applications are equally exacting. A key objective of distribution PMUs is to create a level of situational awareness for the distribution operator that is not available from SCADA (assuming that SCADA monitoring is already in place). PMUs can help detect anomalies or any behavior that does not conform to normal operations, and provide data for applications to issue appropriate alarms to the control center or trigger mitigation strategies. The value of PMU data in this context is substantial, even when the number of PMUs is insufficient to deliver full observability of the grid. Even without complete knowledge, when the state of the

---

[7] It is worth noting that most of the PMU placement literature was developed during the period before PMUs were commercially available at reasonable prices, and most PMU installation entailed a costly and time-consuming process requiring site-specific design, a communications network upgrade, field crew training, dedicated truck roll, and a facility outage to enable the installation. Less than a decade later, in contrast, transmission-level PMUs that meet IEEE technical standards are widely available (many already deployed with functionality resident in many digital relays), reasonably priced, and designs and business practices enable routine installation of PMUs in transmission substations and generation points of interconnection. It should be possible to extend the lessons learned from transmission-level PMU design and installation to save extensive time and money in designing and deploying distribution-level PMUs (whether stand-alone devices or as part of multi-function IEDs). Once low-cost, easy-installation measurement devices and communications are feasible, it will be possible to deploy PMUs across key portions of the grid with limited need to "optimize" device placement.

unobservable part of the network remains ambiguous, it is still possible to more or less coarsely classify the overall state of the system and assess abnormal trends and risks.

For example, the authors in [49] propose a data analytic that is based on the inspection of validity of Ohm's law for the full network, **I** = **YV**,[8] which forces current and voltage phasor measurements (**I,V**) to be orthogonal to (1,-**Y**). These algebraic equations are not exactly met during a fast transient, or exit from the quasi-steady state. This fact can help identify if there is an anomaly – for example, because the grid parameters have changed, or an external transient affected the distribution circuit. Furthermore, this rule gives a theoretical criterion for placing a limited number of PMUs for event detection: the goal is to be maximally sensitive to events that reflect the exit from the steady-state operation of the grid, regardless of the origin of the disturbance. The resulting placement algorithm optimizes a metric that depends only on the grid topology and electrical parameters. Figure 10 illustrates the optimal placement of 20 PMUs on the IEEE-123 test circuit based on this criterion.

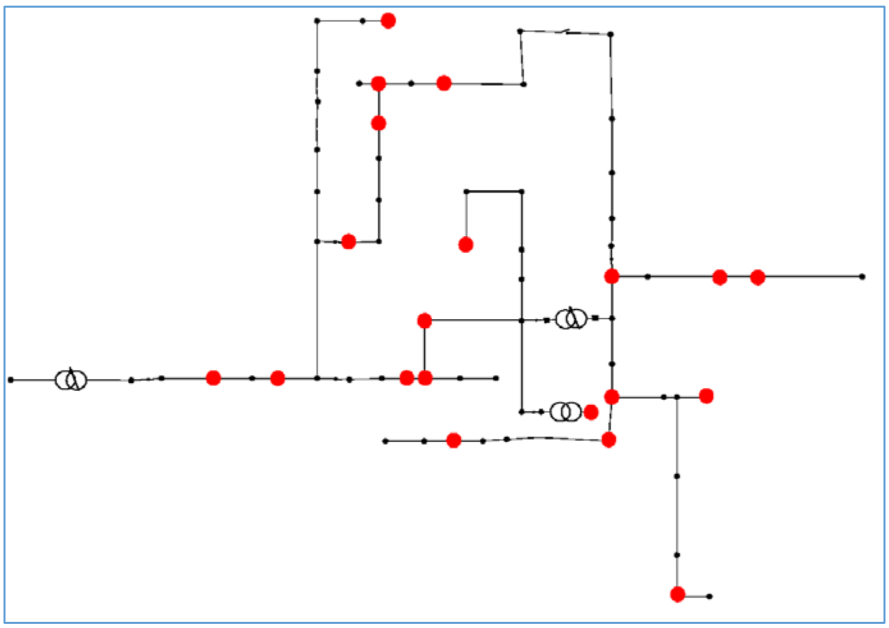

Figure 10: µPMU Placement for Event Detection in IEEE-123 Test Feeder. (ASU)

As Figure 10 shows, the placement strategy scatters the PMUs fairly evenly across the network, and consistent with the intuition that sensors should be placed near the feeder head, the end of the feeder or its laterals, and branch points in the circuit.

While theory can offer useful insights to confirm what is reasonable practice, PMU placement strategy will vary according to the user's priorities and particular circumstances. Distribution circuits are highly diverse in terms of topology and physical layout of facilities, not to mention loads and connected resources, so there is no standard, one-size-fits-all solution for optimal placement. For example, if the goal is to inform the protection of an urban meshed network,

[8] Written here in the succinct form where **I** is the vector of branch currents, **V** is the vector of node voltages, and **Y** is the nodal admittance matrix of the network that reflects the admittances (inverse impedances) connected to each node. Each element within **I**, **V** and **Y** is a complex number or phasor.

this would call for a different PMU placement strategy than monitoring the impacts of high-penetration DER on a long radial feeder. In sum, there is still much to learn about the optimal design of a distribution-specific PMU network.

If, in the future, the cost per installed PMU is small enough to just scatter them liberally across distribution systems, the optimal placement problem may translate into an optimal data usage problem. The question of where to install PMUs would then translate into which PMUs to query for data, for a given application.

4.6 PMU Placement: Practical Aspects

Practical considerations also demand that devices can be installed with reasonable effort. Until PMUs become a standard component for installation alongside or embedded within other circuit devices at the distribution level (e.g. switchgear, capacitor banks, transformers, or inverters), the best locations may simply be those that are most convenient.

The following are observations and lessons learned from early field deployments of µPMUs for the distribution grid [7]:

- It is highly desirable to connect PMUs in locations where a low-voltage signal is already present and accessible, i.e. connecting to the secondary side of existing instrument transformers (transducers) or service transformers. The alternative is to purchase and install new transducers rated for primary distribution voltage, likely at substantial cost. Beyond the substation, there are various possibilities to tap into existing PTs or CTs associated with line devices such as capacitor banks or reclosers, switches, and service panels.
- Outputs from some existing transducers may be incompatible with inputs to µPMU sensors in terms of voltage level, current signal range, or impedance. Additional high-precision 5A:0.333V CTs may be used to adapt the output from typical existing CTs to provide a suitable input for the µPMU.
- Connection to 120V convenience outlets at electrical facilities may be an option, but this offers only single phase measurements.
- Some PMU connection points may be difficult to access, either because they are on customer premises, or because they are surrounded by high-voltage equipment that mandates special safety precautions. Researchers in particular should consider what procedures will need to be followed to access a PMU or its modem in the field. On the other hand, field installations should not be too vulnerable to theft or vandalism.
- Proximity to customers may also trigger data privacy concerns. For example, current measurements on the secondary side of a distribution transformer with fewer than 15 customers connected may be considered a threshold for customer privacy.
- A high-performance GPS receiver requires an unobstructed view of the sky, either immediately from the PMU location, or reachable with cable. Trees and other vegetation and roof structures can present challenges. For indoor installations, placing a GPS receiver at a window may or may not work reliably.

- For data streaming, Ethernet, cellular service, or some other communications infrastructure must be available on location. Cellular modems require a suitable power source.
- The performance of different types of modem antennas with respect to dropped packets and latency varies significantly. There is a trade-off between modem performance and form factor, where a small form factor might be desired for reasons of aesthetics, space constraints, or minimizing draw as targets of vandalism. Some modems may not work well near the service boundary of different signals (3G/4G/LTE).
- Underground locations will have special constraints on access and wireless connectivity.

In sum, logistics and practical considerations are important for determining good and affordable installation sites. Field personnel need to be consulted by design engineers early in the process, and some iteration is to be expected. Since the cost of installation can outweigh the cost of a PMU (on the order of several thousand dollars), it may be less expensive to place more PMUs where convenient, than a minimum set at optimal locations. Such a strategy does impose a greater burden on the data infrastructure, the topic of the next section.

4.7 Communications and Data Collection

A useful PMU network requires local data storage (both resident within the PMU and within a local phasor data concentrator in the field), precise time-alignment and data storage at a central collection point (as in a distribution system control center), analytical tools to use those data, and high-speed, high-performance data communications networks to collect and transport the data.

Synchrophasor data are typically produced at a rate of 30-120 Hz, where one sample per half-cycle is the greatest meaningful time resolution in the phasor domain (any more detail about the voltage or current waveform has to be represented in the time domain). If stored, these measurements quickly add up to a considerable data volume. For example, a device that reports magnitudes and phase angles for 3-phase voltages and currents at 120 Hz produces 124 million data points or about 1 GB per day.[9] Since some applications involve real-time monitoring and others rely on off-line or forensic analysis, synchrophasor databases should serve both needs and support access to both real-time and archival data. With conventional tools, however, searching through archives for specific events at high granularity can take significant time and effort on the part of the user – and that assumes the user knows what they are looking for. This problem could become particularly challenging in distribution systems, where diagnostic applications may refer to numerous, densely deployed PMUs – but where disturbances happen quickly, so that it is generally desirable to preserve all records at full resolution. An effective distribution synchrophasor data infrastructure must therefore support large data volumes (many terabytes), extremely fast searches, and compatibility with advanced automated applications such as machine learning.

[9] 6 phasor values * 2 components per phasor * 120 values/sec * 3600 sec/hr * 24 hr/day = 124 million values/day. Each phasor value and associated time stamp is represented by several bytes.

The synchrophasor data rate is critical for observing fast-changing phenomena such as transient disturbances and faults, but under normal operation, a high data rate carries less information because successive data samples are more correlated. This implies that synchrophasor data can be compressed or decimated as appropriate, before transferring to data processing algorithms at different control layers with different needs.

Since not all applications require the same reporting rate, the frequency of communication can vary – for instance, down-sampled from once per cycle to once every few seconds, or anomaly-triggered. For maximal flexibility, the data should be able to flow to multiple networking nodes, where each node can be armed with different analytic tools. The analyzed data can be visualized at the node, sent to users as is or filtered, or sent when a threshold for anomaly detection is met. For example, a summary may be sent to the distribution system operator, or a control instruction sent from the networking node to the relevant devices [51]. New analytic tools with variable-rate data processing are one approach to reducing the burden of data communications and archival.

4.8 BTrDB

The Berkeley Tree Database (BTrDB) was developed (under an ARPA-E award) specifically for the efficient storage and fast searchability of time-series micro-PMU data in the distribution context [36] [50] [51]. BTrDB is an open-source software that can run on distributed commodity hardware or in the cloud, as part of an innovative architecture for synchrophasor data analysis. It handles large insertion rates while supporting advanced query and visualization techniques for both real-time and historical data through an innovative tree data structure. For example, the database can return search results of 2,000 points summarizing anything from the raw values (9 milliseconds) to 4 billion points (a year) in 100-250 milliseconds [50]. With this tool, human users and automated search tools can seamlessly zoom in and out to scrutinize data streams at different time resolutions, and easily find events buried in terabytes of data like needles in a hay stack. Visualization is provided by a multi-resolution plotter ("Mr. Plotter").[10]

Figure 11 shows the architecture and deployment design for the original BTrDB system, which can scale to handle analyses and storage for tens of thousands of PMUs [51]. This level of scalability becomes necessary in the context of deployments on the order of 10 PMUs per distribution circuit. It also introduces the possibility of comparing synchronized data from thousands of distribution feeders across a wider area, which has never been done before.

As shown in Figure 11, µPMUs stream raw data into the database by way of the "chunk loader." To enable human-centric analyses, the data is then automatically "distilled" using the GPS lock stream and continually evolving heuristics for good data, into globally timestamp-aligned clean data streams. These streams become the inputs for a set of additional algorithms ("distillers") that create a directed data flow graph for a single phase of an individual µPMU. These are

[10] An illustrative sample of µPMU data can be seen at http://powerdata.lbl.gov.

repeated for the other phases, and again for each of the other µPMUs [50]. The BTrDB networking infrastructure is agnostic to the particular sensor device connecting to it, so its use is not limited to PMU data. Rather, it can process and combine time-series data from heterogeneous sources and at different time resolutions. BTrDB employs standard protocols and file formats in order to facilitate a diversity of devices, information, and applications.

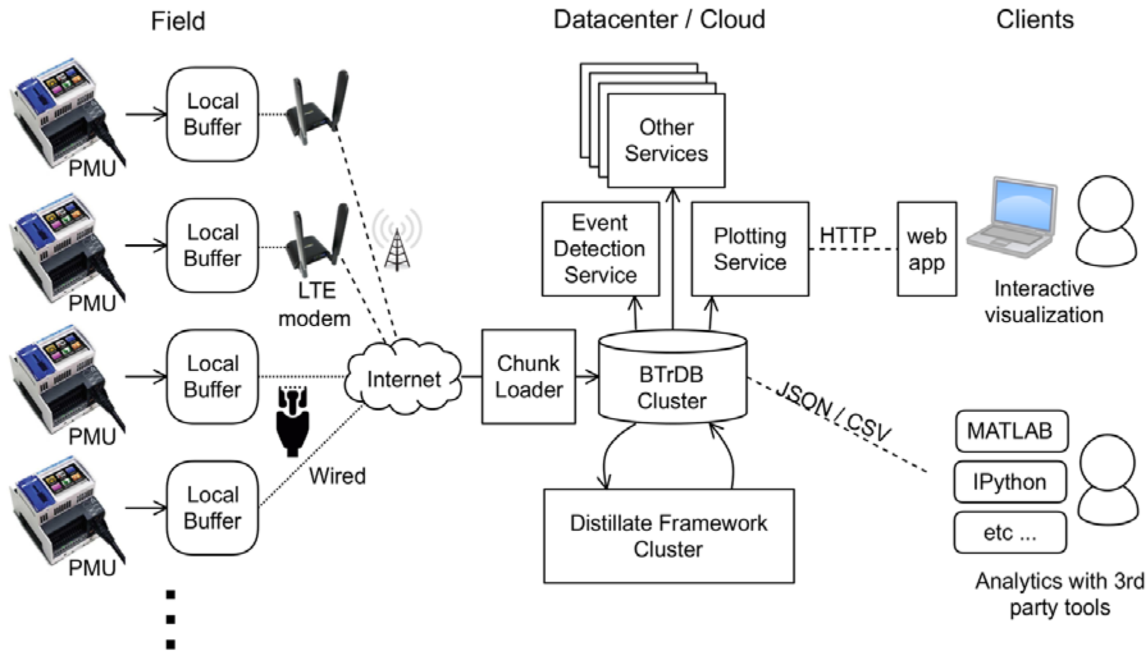

Figure 11: BTrDB Data System Architecture. (Michael Andersen, UC Berkeley)

4.9 Cybersecurity and Synchrophasors

There are two major cybersecurity issues involving synchrophasors that, ultimately, will have to be addressed in utility implementations at the distribution level: (1) to what extent can and should PMU and PDC system installations provide warning and mitigation of cyberattacks; and (2) what minimum defensive measures are essential to offset known vulnerabilities and threats involving synchrophasors.

On (1), we expect that utilities will structure and operate networked synchrophasor deployments to provide integrated situational awareness. Warning tools would include time-synchronized cyberattack anomaly detections, possibly based on techniques like those described later in this report. Mitigation tools may include defensive efforts to maintain electric service, including "islanding" measures as defensive perimeters shrink. As national policies evolve that address critical power dependencies (such as national security installations, air traffic control systems, major medical facilities, etc.), utilities will likely have to build situational awareness capabilities to support survivability in cyber war. Distribution synchrophasor networks could play a unique and vital role in this context.

Regarding (2), it is known that the North American grid has been targeted for potential cyberattacks by foreign entities since at least 2012. In particular, Russian cyber capabilities raise grave concerns on vulnerabilities that extend to the distribution substation level. Known vulnerabilities include "supply chain" vendor penetrations (both grid-specific and IT-industry components), "human-machine interfaces" (HMI) for control, "industrial control systems" (ICS) including firmware replacement, multiple methods for penetration of networks and connectivity protocols, vendor-maintained systems, and attack structures for remote control, stealth and anonymity. Critical Infrastructure Protection (CIP) v5/v6 standards have proven ineffective to detect or mitigate recurring Russian grid reconnaissance, which now include nuclear generation sites. While specific recommendations are beyond the scope of this report, we note the importance of designing distribution synchrophasor networks with an ability to resist cyber-intrusions.

4.10 Time Synchronization Requirements

For correct operation at the required performance level, all elements of a synchrophasor system (both PMUs and the associated phasor data concentrators) must continually access a common and accurate timing source linked to Coordinated Universal Time (UTC). Any inaccuracy in a PMU's timing adversely affects the PMU measurements, especially the estimation of phase angles of the measured quantity.

The IEEE C37.118.1-2011 standard suggests a maximum uncertainty in the synchrophasor time stamp of 1 μs. This standard can be met by referring to a conventional GPS pulse-per-second signal, or some other high-accuracy time delivery method. Alternative or supplementary high-precision timekeeping technologies such as eLORAN, White Rabbit, and chip-scale atomic clocks have been discussed in the synchrophasor context [52]. GPS receivers or network-distributed time are presently the most common methods for accessing precision time signals.

The 1-µs accuracy standard for timekeeping recognizes that there will be other sources of error besides the reference clock, including the magnitude measurement, which all need to fit into the error budget (equivalent to 26.5 µs at 1% TVE). Another important source of error originates from outside the PMU proper, namely, the voltage or current transformers through which PMUs are connected to medium- or high-voltage equipment; these transducer errors are not part of the PMU performance standard, but are discussed in a separate section below.

Recent testing performed on commercial PMUs (documented in [53] and using tools found in [54]), supports the case for strict time accuracy standards, leaving more margin for other errors. Figure 12 illustrates the effect on phase angle error of an artificially introduced time error. The first step change with a 10-µs error already drives the phase angle error near or above the total PMU error budget, shown as the reference line at 0.573°. Further discussion can be found in the Singh et al. paper on assessment [55].

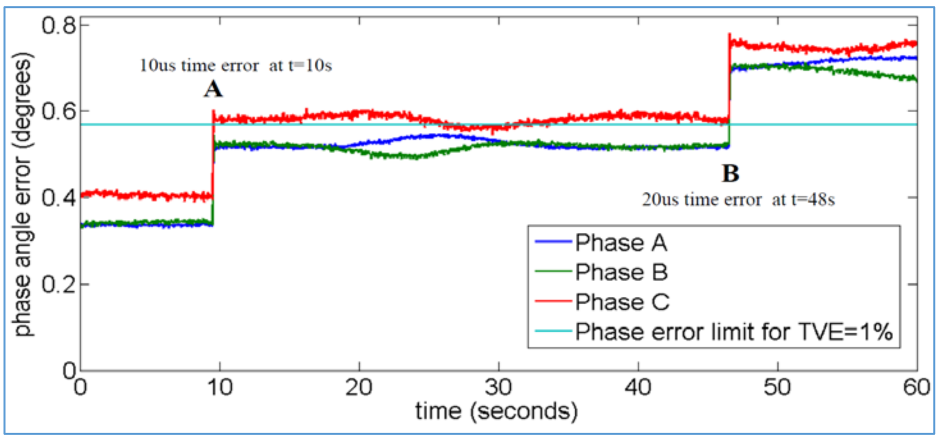

Figure 12: Phase error in measured voltage signals in presence of instrumentation channel and varying time error (leading time) in steps of 10µs at point A and B. [53]

In the context of distribution systems, there is reason to consider synchrophasors with greater precision and accuracy than required by the C37 standard. How best to articulate performance expectations or future standards for this context will be an important subject of further work. It may turn out that TVE is not the most suitable way to describe a phasor measurement error in situations where the acceptable error is very small, e.g. where measurements from different locations are very similar to each other. Appendix A offers some comments intended to stimulate discussion in this direction.

Appendix A: Synchrophasor Representation and Quantities

A.1 Phasor Representation

The term "phasor" implies an abstract and idealized mathematical representation of alternating current (a.c.) electrical quantities. It takes the form of a vector with a magnitude and angle, where the magnitude represents the amplitude or root-mean-square (rms) value of voltage or current, and the angle represents a time shift in the waveform. Though it is an essential tool in power engineering, phasor analysis becomes problematic if and when the actual power system departs from the idealized state. In fact, how best to define phasor quantities and their accuracy under dynamic conditions is not a settled question.

Section 2.0 provided a brief summary of key terms that express the conventional, idealized understanding of synchrophasor measurements and errors. For many uses of synchrophasors, that understanding is quite sufficient. But with the introduction of distribution synchrophasor applications comes a new level of expectation for the performance of PMUs, and a new round of questions about the fidelity with which phasor quantities represent what it physically happening in the electrical network. To help support these conversations, this Appendix will walk through the details of the conventional definition, and then discuss some of the nuances and theoretical limitations of synchrophasor measurements.

It is convenient and standard to assume voltage and currents in an a.c. power system are sinusoidal functions of time. Written in cosine form, the equation for a single such quantity is

$$x(t) = X_m \cos(\omega t + \phi) \qquad (1)$$

Here $x(t)$ is the instantaneous value of the function, X_m is the maximum value, ω is the angular frequency (2π times the a.c. frequency in hertz) and ϕ is usually called the phase. This simple wave representation has given rise to a number of very useful methods in power systems. Most useful among them is the phasor diagram.

Power system engineers commonly assume that the frequency is at its nominal value, and refer to just the magnitude, X_{m}, and phase, ϕ, as a *phasor.* These two quantities, as measured and time-stamped by a PMU, are referred to as a *synchrophasor.* By convention, power engineers replace the maximum value or amplitude X_m with the root-mean-square value, $X_{\mathrm{rms}} = \frac{1}{\sqrt{2}} X_m$. The local frequency is also measured and reported by the PMU.[11] For the purposes of measuring ϕ, the time may be taken as a "local zero" at the UTC second rollover, and therefore at the top of every cycle of the reference signal.

Often, a different representation is used, based on Euler's equation.[12] For a complex signal $Ae^{j(\nu t+\varphi)}$ the relationship

[11] While the frequency of a synchronous a.c. system is the same everywhere to a first approximation, precise measurements reveal local and short-term departures from synchronicity, which are of particular interest here. For example, power oscillations across an a.c. grid can be observed in terms of changes in local frequency.

[12] $e^{jx} = \cos x + j \sin x$, where $j = \sqrt{-1}$.

$$Ae^{j(\nu t+\phi)} = A\cos(\nu t+\varphi) + jA\sin(\nu t+\phi) \quad (2)$$

allows us to regard the projection of the sinusoidal quantity onto the horizontal axis as being represented by the real part of the quantity.[13] That is, an equation of the form of Equation (1) is given by

$$A\cos(\nu t+\varphi) = \mathrm{Re}[Ae^{j(\nu t+\phi)}] \quad (3)$$

Likewise, we may write,

$$X_{\mathrm{m}} cos(\omega t+\phi) = \mathrm{Re}\left[X_{\mathrm{m}} e^{j(\omega t+\phi)}\right]. \quad (4)$$

We can separate the exponential expression into two parts:

$X_{\mathrm{m}} e^{j(\omega t+\phi)}$ can be written $X_{\mathrm{m}} e^{j(\phi)} e^{j(\omega t)}$.

The time-dependent, unit-sized part $e^{j(\omega t)}$ is called by mathematicians the *rotating phasor*, or the *rotator* [56] [57] [58].[14]

The part $X_{\mathrm{m}} e^{j\phi}$is called the *stationary* phasor. We could write that stationary phasor as $X_{\mathrm{m}} e^{j\phi} = \boldsymbol{X}_{\mathrm{m}}$ so that the complete expression $X_{\mathrm{m}} e^{j(\omega t+\phi)}$ is $\boldsymbol{X}_{\mathrm{m}} e^{j\omega t}$.

The complex number $\boldsymbol{X}_{\mathrm{m}}$ gives the position of the rotating phasor at time *t* = 0. This corresponds to the part of Equation (4) (or equivalently, Equation (1)) that power engineers call the *synchrophasor*.

The phasor representation brings significant simplification to the problem of analyzing power systems, because it allows the use of geometrical solutions to the equation representing the signal, instead of the more complicated trigonometric ones. Further, it allows the use of what are called phasor diagrams, an invaluable aid to visualization. A phasor diagram representing Equation (1) is shown on the right side of Figure 13. The left side shows how it is derived. Some power engineers do not draw the arrow indicating rotation, but it serves as a useful reminder that the object we treat as stationary for the purpose of doing arithmetic is imagined as spinning counterclockwise at the (presumed constant) grid frequency.

---

[13] A sinusoid as the real part of a complex function can be visualized as the shadow cast by a three-dimensional helix or corkscrew, which appears as a sine or cosine wave when viewed from any side (perpendicular to its central axis, which in our case represents time).

[14] Recall that the quantity e^{jx} rotates in the complex plane as x increases, and $|e^{jx}| = 1$ for all x.

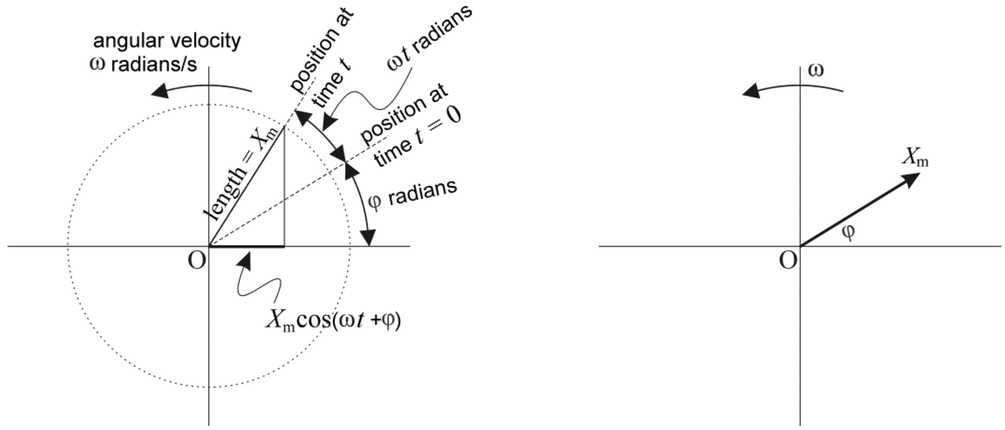

Figure 13: Rotating phasor whose projection
on the horizontal axis is $x(t) = X_m\, cos(\omega t + \phi)$. *(PNNL)*

Note that the word "phase" can be interpreted as the entire argument of the cosine ($\omega t + \phi$), or as just the value at time-zero, essentially the phase offset (ϕ). In the PMU, phase is the value measured as the phase difference between two signals: one is the signal being observed, the other is a reference, an assumed sinusoid at exactly the nominal frequency (50.00 or 60.00 Hz) timed such that the positive peak of the signal coincides with the second-tick of UTC. Phase for the PMU is thus measured between two signals that may be of different frequency. Therefore, the difference in the values of the entire argument of the cosine must be used. Time can be taken as zero at the start of any cycle of the reference wave.

In summary, a PMU reports the values of a *synchrophasor*, namely, the result of the measurement[15] of the magnitude, X_{m}, and phase angle, ϕ, relative to a cosine function at the nominal power frequency. Phasor values are the values of an abstraction known as the "analytic signal," and the values found by measurement apply for the entire duration of the observation window (say, one a.c. cycle). The analytic signal is discussed further in Section 4.

Figure 14 illustrates a typical evolution of phase angle reported by a PMU as the frequency increases from below 60 Hz to a frequency higher than that. When frequency happens to be different than 60.00 Hz, the angle relative to the reference steadily changes, and periodically "wraps" around between 360° and 0°. Absent local disturbances, PMUs at two locations in the power system will show a constant difference between their phase angles, and their traces appear parallel except that they wrap at slightly different times.

---

[15] It is worth remembering that the word "measurement" refers to a process; the result (say, a measured value) is what we get out of the process.

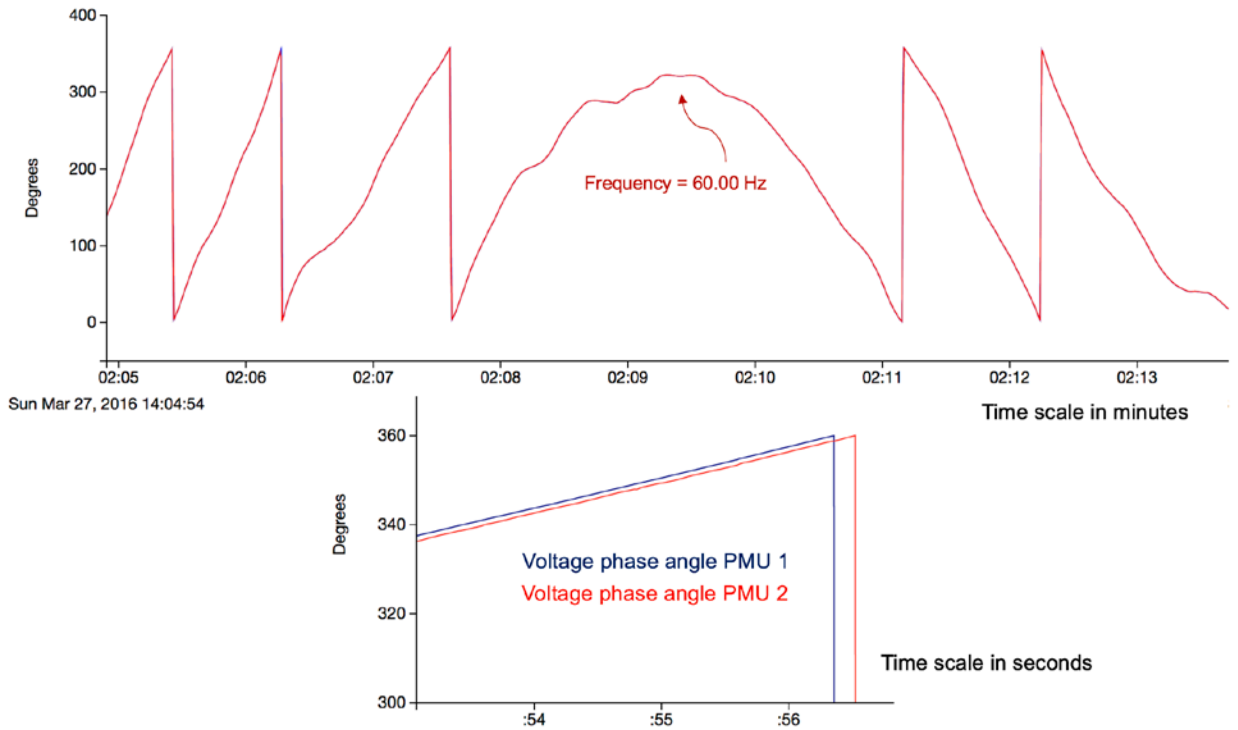

Figure 14: Phase angles reported by PMUs. (UC Berkeley)

A PMU generally also reports a frequency ω. Note that if ω were in fact exactly constant (as presumed by the definition of a 'stationary phasor'), and at the nominal value, the phase angle would never change. If the frequency is constant but not at the nominal value, the phase will generate a ramp with time, as in the left and right sides of Figure 14. Only when the grid is at the nominal frequency will the phase value be constant.

In practice, it is understood that grid frequency varies – but hopefully only slightly. If the change in ω is relatively small and slow, the time rate of change of ϕ in the stationary phasor reported by a PMU is not particularly informative. It is also understood, however, that if the measured value of ω changes very suddenly – in other words, if the actual electrical quantities seriously fail to conform to the assumed sinusoidal model – there arises an ambiguity as to what a PMU should report as the "phase" and the "frequency". This sort of event can happen during a fault on the power system. It is worth remembering, though, that PMUs were not designed to cope with conditions so far from normal.

The quantities X_m, ϕ and ω are evaluated over an interval of time defined by the IEEE standard C37.118.1 [59]. Because of the inherent ambiguity, this standard does not require a rigorous reconciliation during transient conditions. This ambiguity is further discussed in Appendix B.

The rate of change of frequency (ROCOF), expressed in hertz per second, describes how rapidly the frequency is changing, indicating an imbalance between generation and load. It is represented by a term in the cosine argument that varies as the square of the time. In the ideal steady state, ROCOF would be zero.

A.2 Total Vector Error

The time accuracy requirement in IEEE C37.118.1-2011 is indirectly determined by the need to meet the requirement for a maximum 1 percent of a value called the Total Vector Error (TVE), under steady-state conditions. TVE is a way to express the uncertainties in the result of the measurement as a combined error budget for two components: one due to the measurement of amplitude, and the other due to the measurement of the phase. In the case of zero uncertainty in the result for phase, a maximum of 1% or 0.01 per-unit error is allowed in the result for amplitude, and vice versa. A problem with TVE is that because it combines two quantities, their respective uncertainties cannot be propagated in a rigorous manner.

The TVE is easy to visualize as a circle around the tip of the synchrophasor arrow, where the radius of the 'error circle' is simply 1% of the full length of the phasor being measured. This is illustrated in Figure 14.

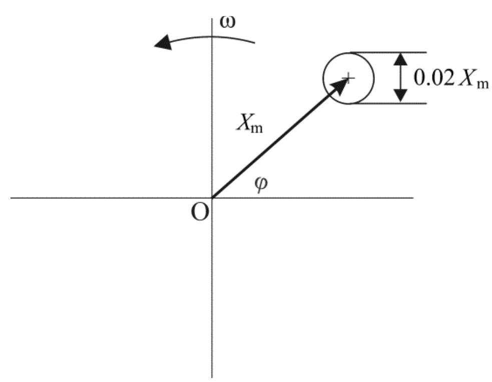

Figure 14: Visualization of Total Vector Error, not to scale. (PNNL)

It is not obvious how to put a percentage error number on phase, because phase is somehow "different" from magnitude. The quantities that are standardized by the SI system are what are now called *rational* in the classification scheme of Stevens [60] [61] [62]. The term implies something about the kind of mathematical operation that can be done on the result of the measurement: it is possible to take *ratios*, because all the SI quantities have a single natural zero and a linear scale. Angle is a fundamentally different kind of quantity. The quantity ϕ has to be evaluated from some defined reference, and should really be called a phase *offset.*

The allowable uncertainty in phase is determined as follows. To cope with the difference in kind, it may be conjectured that at some point in the history of the PMU it was decided that if the 1% uncertainty of amplitude were regarded as a small phasor added to (or subtracted from) the phasor that represents the power system quantity being measured, angle uncertainty and amplitude uncertainty could be combined, much as the uncertainties in rational quantities were combined: as the square root of the sum of the squares.[16]

---

[16] This method of combining is actually justifiable only under a restricted set of conditions that were not taken into consideration in this development. Nevertheless, the method seems to have produced a measure of uncertainty that applies to angle with a size that seems quite appropriate to our situation.

For two variables, this combination method, the root of the sum of the squares, results in the equation of a circle. Recall that the equation of a circle whose center is the origin is

$$r^2 = x^2 + y^2. \tag{5}$$

A circle whose center is (*a*, *b*) is described by

$$r^2 = (x-a)^2 + (y-b)^2. \tag{6}$$

Added to a phasor diagram, the terms *a* and *b* are the horizontal and vertical offsets of the center of the circle, so that to put the circle at the end of the phasor, Equation (6) becomes

$$r^2 = (x - X_\mathrm{m}\cos(\omega t + \varphi))^2 + (y - X_\mathrm{m}\sin(\omega t + \varphi))^2. \tag{7}$$

The circle thus described is not a phasor, though it is sometimes drawn as if it were in order to illustrate the size of the circle. It is simply a circle that moves along with the tip of the phasor, and there describes a region within which the total vector error is defined as being acceptable.

This situation is illustrated in Figure 14. For clarity, we further exaggerate the size of the TVE circle in Figure 15.

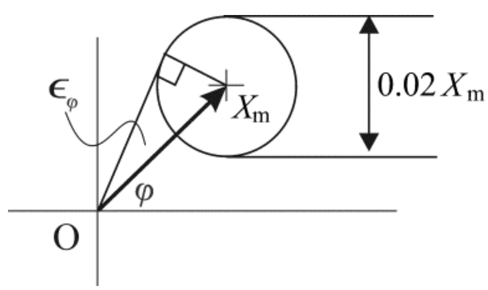

Figure 15: Close-up visualization of Total Vector Error, very much not to scale. (PNNL)

Here we can see that the geometry of the situation indicates that the maximum permissible phase error is given by

$$\epsilon_\varphi = \sin^{-1}\left(\frac{1}{100}\right) = 0.573° \tag{8}$$

where ϵ_φ is the phase error. The fact that this value is a constant supports the earlier observation that angle is somehow “different” from magnitude. A 1% error in magnitude depends on the size of the quantity being measured. A “1% error in angle” is always just over half a degree, whatever the value of the phase being measured.[17]

While the origin of the TVE presented above is speculative, the result is not. For the case of zero magnitude error, where the entire 1% error budget can be allocated to the angle portion, the maximum permissible phase angle error within the 1% TVE standard is 0.573°. This error is independent of the angle being measured.

At 60 Hz, where a full period of 360° corresponds to 1/60 of a second or 16.7 milliseconds (ms), an angle error of 0.573° corresponds to a timing error of 26.5 microseconds (μs). If the clock signal is delayed, the phase will be reported as advanced from its proper value. The standard suggests a maximum timing uncertainty of 1 μs, rather than 26.5, to allow for sources of uncertainty other than angle.

It is interesting to examine the way phase and magnitude errors combine. Figure 16 shows a situation where the measured angle and magnitude each have errors.

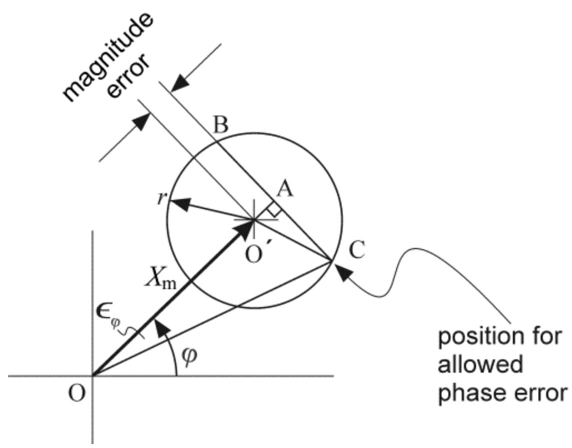

Figure 16: Combined magnitude and phase error. (PNNL)

In Figure 16, the center of the TVE circle, radius *r*, has been labeled O´. Suppose that the phasor magnitude OO´ is measured as OA, larger than it really is. A line has been drawn from O to A, passing through O´. A chord of the circle has been drawn perpendicular to the line OA. It intersects the circle at B and C. The maximum allowable phase error therefore corresponds to one of the two points at which the chord BC intersects the circle, say at C. The problem is to find the angle ϵ_φ.

---

[17] The value, just over 34 minutes of arc, seems to pass a “reasonableness test.” The IEEE Standard for instrument transformers (C57.13-2008) has three classes of accuracy. They allow a magnitude error of 0.3, 0.6 and 1.2 percent. The maximum angle error corresponding to each of these classes is 30, 60 and 120 minutes. One might reasonably infer that a 1% transformer would be allowed 100 minutes of phase error, just three times the amount allowed in the PMU.

The triangle O´AC is a right triangle with hypotenuse *r*. The angle O´AC is a right angle by definition. The line O´C is of length *r*. If we denote the length O´A by ϵ_{m}, then the length of the line AC is given by $\mathrm{AC} = \sqrt{r^2 - \epsilon_{\mathrm{m}}{}^2}$. It follows that the angle ϵ_φ is given by

$$\epsilon_\varphi = \sin\left(\sqrt{r^2 - \epsilon_{\mathrm{m}}{}^2}\right).$$

This straightforward Pythagorean relationship allows us to draw graphs showing the interaction of the two sources of error. Examples (similar to the examples in Appendix E of the Standard) are shown in Figure 17.

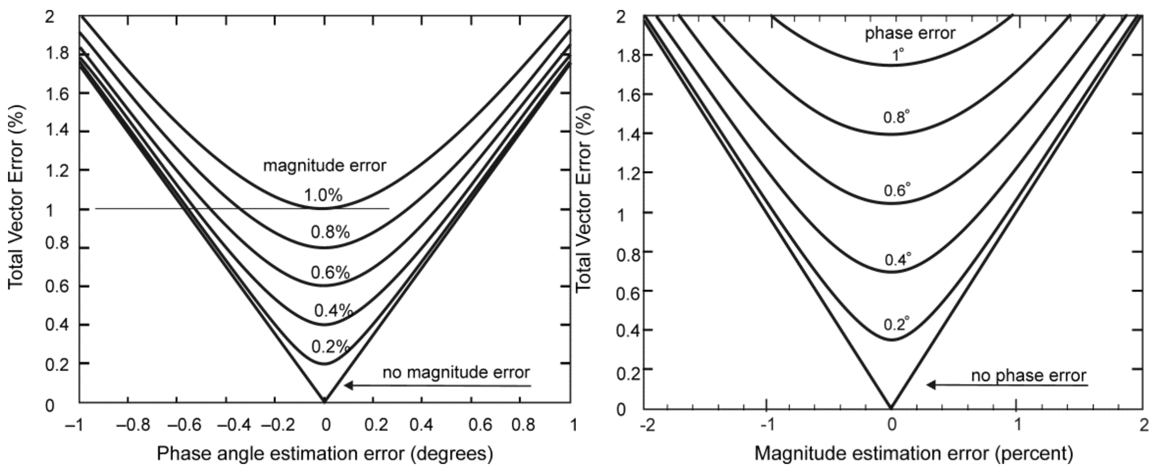

Figure 17: Interaction of phase and magnitude errors. (Based on Figs. E.2 and E.3 of IEEE Std C37.118.1-2011)

A.3 Comment on TVE

The TVE definition has the solved a problem for the PMU: it allows a “percentage error” number to be attached to the measurement of phase. Unlike the result of a measurement of a rational quantity, this percentage error is constant, regardless of the size of the angle being measured.

However, the combining of two separate measurement processes within the TVE may also be counterproductive. Measurement results are normally stated with a pair of numbers giving the uncertainty of the result. Because of TVE, it is impossible to do that for the angle or the magnitude measured by the PMU.

A potentially more useful approach would be for the PMU community to accept the value of 0.573 degrees as corresponding to 1%, but also begin to state the measurement results with a two distinct statements of uncertainty: one for angle and another for magnitude.

Such a change would bring the PMU into line with the way IEC standards specify accuracy requirements for instrument transformers in IEC 60044-1 [63]. For example, a measuring CT of accuracy class 0.2 would be allowed an error of no more than 0.2% ratio at rated current, and a phase displacement of no more than 10 minutes of arc. The limits of error for the complete set of classes in IEC 60044-1 are shown in Figure 18, along with a line defining the 1% TVE for a PMU.

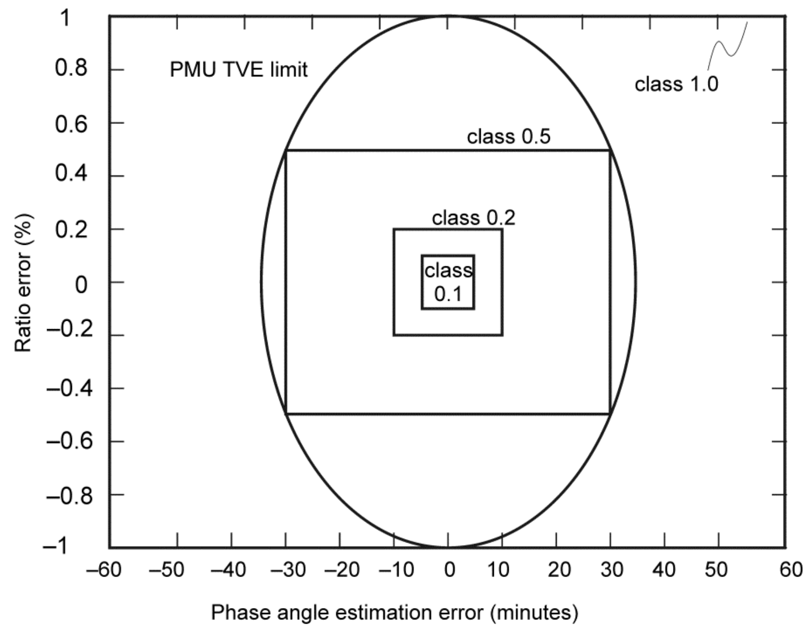

Figure 18: IEC error limits for measurement transformers, along with PMU TVE limit. (PNNL)

Diagrams such as Figure 18 show the limits of accuracy for two parameters. In the IEC standard, they are not independent: although a device may be compliant when it has the maximum allowable ratio error and at the same time the maximum phase error, the values of these limits are related. It is not possible to obtain a device with only 30 minutes of phase error without at the same time specifying a maximum 0.5% ratio error. In the sense of connecting ratio error and phase error, TVE is no different.

However, in the sense of being able to propagate uncertainties, TVE *is* different. TVE has only one class, and does not allow the user to know the separate errors on phase and magnitude. An application such as a distribution system voltage control must know the voltage to much greater accuracy than 1%, and an application that computes steady-state power flow along a distribution feeder must know the phase angle difference between two locations to much better accuracy than 0.573 degrees. While that may well be possible with a given PMU, it is impossible to be certain if its performance is described only by a total vector error.

It would be better for some users to define the allowable errors separately.

A.4 Phasor Quantities – A Closer Look

The *stationary phasor* is often called simply the *phasor* by power engineers. The expression $X_{\mathrm{m}}e^{j(\omega t+\phi)}$, evaluated at time *t* = 0, results in the two quantities that define the stationary phasor, X_{m} and ϕ. Power engineers are taught that since the power system frequency is constant, power system signals can be represented by just these two quantities, amplitude and phase. In reality, the signal observed on a power system is not a perfect sinusoid of constant frequency. This section discusses how a phasor can nevertheless be defined mathematically. It is worth remembering, though, that since a phasor always aims at describing a cosine, regardless of how it is obtained, it will not capture information about the signal that is not well modeled by a cosine expression.

Having said that, it is often convenient to assume that the sinusoidal representation applies to voltages and currents in the power system. It is also convenient, in some PMUs and in the analysis of distortion, to perform Fourier transforms on the signals thus represented. If that is done, the spectrum produced is two sided, symmetrical around zero frequency.[18] It may seem troublesome to deal with a negative frequency, since it has no physical meaning. But in fact the negative part of the spectrum is redundant – it contains no information that is not also in the positive part – and it can be removed mathematically.

Figure 19 illustrates the symmetric amplitude and anti-symmetric phase of the power spectral density of an arbitrary real signal $x(t)$.

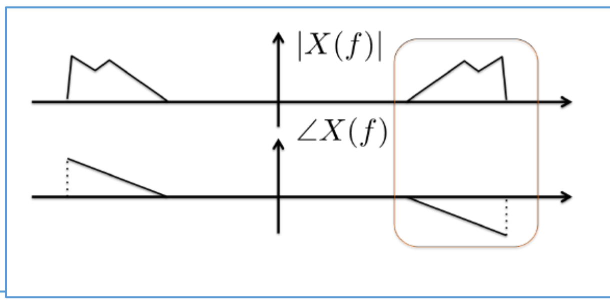

Figure 19: Spectrum of a real valued signal. (ASU)

We assume that the signal of interest is a real-valued function: that is, its values as a function of time are in the set of real numbers. This follows from an assumption made in the process of measurement, namely that the results of the measurement of a series of sampled values map into the set of real numbers. It is intuitive that if the samples represent analog physical quantities such as instantaneous voltages and currents, they must be real.

[18] The frequency spectrum describes a decomposition of the original signal, which can have any shape, into contributions from pure sinusoidal waves at (perhaps infinitely) many different frequencies. The spectrum shows the amount of each that goes into the "recipe" for creating the shape of the composite signal. A helpful introductory reference on this topic is *The Intuitive Guide to Fourier Analysis & Spectral Estimation* by Charan Langton and Victor Levin, Mountcastle Academic, 2017.

Given a real-valued function $s(t)$, we can construct its *analytic signal* $s_a(t)$*:*

$$s_a(t) \triangleq s(t) + j\hat{s}(t) \tag{5}$$

where $\hat{s}(t)$ is the Hilbert transform[19] of the signal $s(t)$.

The analytic signal representation is capable of dealing with time-variant magnitude and angles; in ways that some might consider simpler than the sinusoidal representation. Written in polar coordinates, we could imagine a generalized signal that has a changing phase as $s_a(t) = Xe^{j\phi(t)}$ where $\phi(t)$ is the instantaneous phase. In power system terms, this instantaneous phase, changing with time, could include changes due to changes in both machine speed and system parameters. We can say that the time derivative of phase is the instantaneous (angular) frequency.

This much is hinted at in the IEEE standard for PMUs, but the words must be interpreted with caution. The words "frequency is the derivative of phase" do not apply to the equation $x(t) = X_{\mathrm{m}} \cos(\omega t + \phi)$ if "phase" is taken to mean only the term ϕ. Rather, the "phase" whose rate of change matters is the entire argument of the cosine, $(\omega t + \phi)$.

In reality, neither amplitude nor phase nor frequency of the signal we deal with in the power system are truly represented by constant values. They are themselves functions of time, giving rise to the complex signal $\boldsymbol{X_m}(t) = X_m(t)\, e^{j\phi(t)}$**.** This fact does not preclude the use of phasors, but it suggests constant vigilance to remember the distinction between idealized mathematical objects and the physical phenomena they are intended to describe.

Generally, *any arbitrary real signal* $x(t)$ admits representation into a phasor with respect to *an arbitrary* frequency ω. To understand that, one has to think about what the phasor signal represents in the Fourier domain.

Any real signal $x(t)$ has a power spectral density that has conjugate symmetry around frequency zero, as shown in Figure 19. [20] The phasor signal spectrum is what is seen isolated in the box re-centered on (positive) frequency. In fact, the phasor representation is obtained mathematically by multiplying by zero everything that is at negative frequencies, obtaining what is called the *analytic signal* $x^+(t)$. To obtain the *stationary phasor* (the thing most power engineers recognize as simply the *phasor*) one then shifts the frequency spectrum of the analytic signal towards the origin by ω*.* In the time domain, this implies the relationship $\boldsymbol{X_m}(t) = x^+(t)\, e^{-j\omega t}$. If the (constant) frequency term ω for constructing the phasor quantity is chosen to be around the center of the non-zero support of the power spectral density of the original signal (in other words, if ω characterizes the signal nicely), then the power spectral density of the phasor will have the same profile as that highlighted in the box in Figure 19, but centered around the origin.

[19] A Hilbert transform is a mathematical process that effectively rotates a function in the complex plane, which can be used to make negative frequencies positive.

[20] In other words, the recipe contains the same amount of positive and negative contributions at each frequency.

In other words, it is possible in principle to produce the analytic signal, or express a function as a phasor, for any function (as long as it is real-valued) and any frequency. But the operation that maps an arbitrary signal $x(t)$ onto its analytic signal $x^{+}(t)$ requires the implementation of the so called Hilbert filter, whose response is non-causal and infinitely long.[21]

This calls for an important modification relative to what an actual PMU might do. Since a Hilbert filter is non-causal, it cannot be applied to an incoming signal in real-time, because the signal's shape in the future is not known yet.

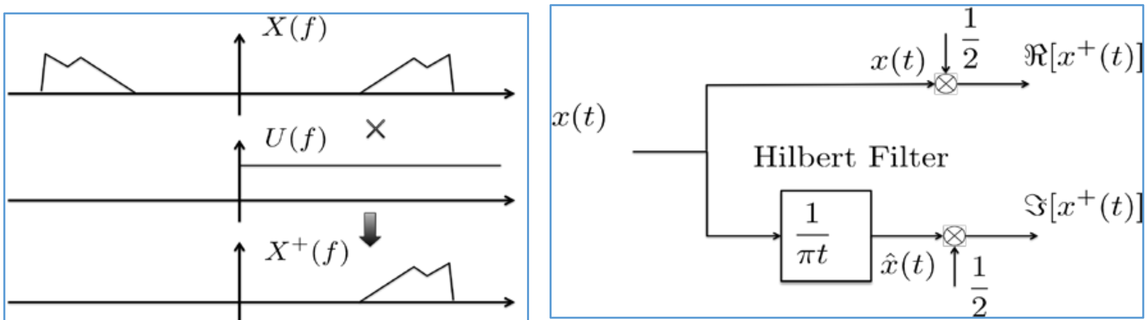

Figure 20: Application of the Hilbert Filter. On the left side of the diagram, the analytic signal is obtained by multiplying the spectrum of the original signal by zero at all negative frequencies and one at all positive frequencies. This operation in the time domain is equivalent to the process shown on the right side. Here the real part of the analytic signal is equal to the signal itself divided by two, while the imaginary part is obtained by filtering the signal through the response 1/ πt. This is called the Hilbert filter. (ASU)

In practice, the analytic signal can still be obtained by what we may call a real (as opposed to an ideal) vector analyzer. Instead of a Hilbert filter, a band-pass filter can be used. This makes a presumption about the original signal: that it has a narrow frequency support around the frequency ω (in other words, it should contain mostly frequencies very similar *to* ω). In effect, we are trading some generality (applicability to arbitrary signals) for the ability to filter the signal in real-time, using only information collected through a specific time window up to the present.[22]

[21] A non-causal filter is one that depends on future inputs, not just the past (which would be causal). The Hilbert filter is "omniscient" in the sense of using information from everywhere along the time axis.

[22] The problem of sampling a signal during a finite time window – what statements can be made about the signal based on a suitably brief observation – is not trivial. A related but distinct problem is the trade-off between latency (speed of reporting) and accuracy of a PMU: given a long time to observe the signal and perform computations, it is easier to produce a nice result. These issues will not be further discussed here.

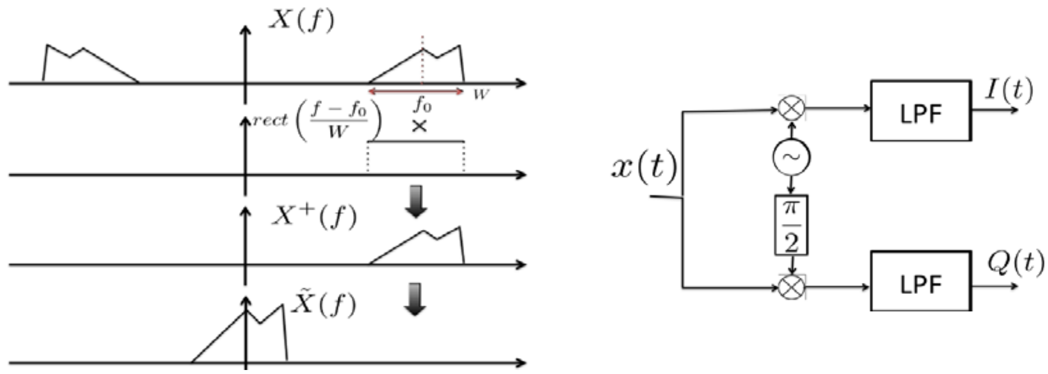

Figure 21: Application of a band-pass filter. When the power spectral density of the signal is narrowband around the frequency ω=2π f_o the extraction of the phasor from the original signal can be done through a band-pass filter instead of using the Hilbert filter. The right diagram illustrates mixing the signal with two in-quadrature carriers and then using a low-pass filter (LPF) to produce the real and imaginary parts of the phasor, $I(t)$ and $Q(t)$, respectively. The C37.118.1 standard includes two different LPFs, referred to as the P and M filters. (ASU)

The practical implementation follows an architecture like the one on the right side of Figure 21, which is similar to Figure C.1 of the C37.118.1 standard. The signal is mixed with two carriers in quadrature ($\pi/2$ apart), namely $\cos(\omega t)$ and $-\sin(\omega t)$. The output of the two mixers is sent to two low-pass filters (LPF). The LPFs extract the real and imaginary parts of $\boldsymbol{X}_m(t)$, respectively: $I(t) = X_m(t)\cos[\phi(t)]$ and $Q(t) = X_m(t)\sin[\phi(t)]$.

The filtering process described here assumes the signal is non-zero only around ω in the original signal. This model implies that we can ignore contributions to the waveform from any frequencies very different from the fundamental ω (such as harmonics), and also that the fundamental frequency changes only slowly. If and when the frequency changes rapidly, the representation of the original signal in terms of a phasor with a certain assumed frequency ω becomes problematic.

Several difficulties might be encountered as a result of a mismatch between the ideal sinusoidal model and the real-world signal. For one, there are likely to be harmonics on the signal, and there may be inter-harmonics. Another challenge is the accurate expression of Ohm's law. Written in phasor quantities, Ohm's Law $I = YV$ only applies to the ideal steady state. While voltage and current phasors are still defined in the non-steady state, they may not obey the expected relationship. In general, we must consider a circuit's response using convolution, $i(t) = y(t) * v(t)$. If the input phasor (voltage or current) changes smoothly, then its relationship with the output phasor (current or voltage) can still be expressed as an algebraic equation, with an impedance or admittance that can be evaluated at the center frequency ω. But if there are sudden changes in voltage or current, the relationship between these variables is not memoryless (i.e. it depends on their history), and even the notion of a time-varying impedance can become inadequate.

In conclusion, the mathematical process of extracting phasors from a signal will produce results in any case. It is possible to define a unique phasor to represent a signal, regardless of how well this signal conforms to our assumptions. This means that it is possible to stipulate a theoretically desired output which PMUs should strive to match in practice, even during non-steady state conditions. But this desired output, the calculated phasor, is itself an abstraction, not a real thing. The fact that the mathematical process is well-defined implies nothing about whether the resulting phasor happens to be an especially useful abstraction. The results will be informative about the real world only to the extent that the real world matches the model assumed while creating the abstraction.

It has been proposed that the phasor measurement problem can be understood as a curve fitting problem, in which three parameters (amplitude, frequency and phase) of the presumed sine wave are to be estimated from the available sample values.[64] [65]. This viewpoint supports the explicit definition of a “goodness of fit” measure, which quantifies how well the best sinusoidal characterization actually matches the original signal. PMUs – any PMUs, based on any measurement method – could determine and report the “goodness of fit” as a real-time quantification on the fit of the implicit “phasor” model.

Appendix B: Public Distribution Synchrophasor Efforts and Deployments

This Appendix reflects the authors' best knowledge at the time of writing and makes no claim of completeness. Note that it does not include information about private projects.

UNITED STATES			
Years	Funder	Project	Team
2003 - present	**DOE**	**Frequency Monitoring Network (FNET/GridEye)**	The University of Tennessee, Knoxville, Oak Ridge National Laboratory.
Objective: Develop a low-cost, quickly deployable GPS-synchronized wide-area frequency measurement network. This is a pilot WAMS and the first one in distribution level, with over 270 FDRs (the single-phase synchrophasor for FNET/GridEye) deployed. Online applications include power system disturbance detection and location, oscillation detection and analysis, islanding detection, frequency and phase angle visualization, ambient data based modal analysis, FIVDR detection, and harmonics monitoring. There are also several offline applications such as model validation, post event analysis, and data analytics.			
Links: http://fnetpublic.utk.edu/, http://powerit.utk.edu/			
2013 - 2016	**DOE ARPA-E**	**Micro-Synchrophasors for Distribution Systems**	CIEE, UC Berkeley, LBNL, Power Standards Lab with utilities Riverside PU, SCE, Southern Co. and others.
Objective: Develop & test micro-PMU hardware and software, field deployments, explore distribution applications, feasibility and data requirements.			
Link: https://arpa-e.energy.gov/sites/default/files/5_CIEE_von%20Meier_GENI3.pdf			
2015 - 2016	*** Work came out of above project**	**Open Micro-PMU LBNL Public Dataset (www.powerdata.lbl.gov)**	LBNL.
Objective: Produce and make publicly available a real-world micro-PMU data set for research and application development.			
Link: https://github.com/lbnl-cybersecurity/lbnl-cybersecurity.github.io/blob/master/LBNL-1006408-Open-uPMU_IEEE_Oct2016.pdf			
2015 - 2017	**DOE**	**Microgrid Research, Development, and System Design**	ALSTOM/ General Electric.
Objective: Research, development and testing of advanced commercial-grade microgrid controllers.			
2015 - 2018	**DOE CEDS**	**Supporting Cyber Security of Power Distribution Systems by Detecting Differences Between Real-time Micro-Synchrophasor Measurements and Cyber-Reported SCADA**	LBNL, ASU, EPRI, EnerNex, PSL.
Objective: Design and implement a measurement network, which can detect and report the resultant impact of cyber security attacks on the distribution system network.			
Link: http://crd.lbl.gov/departments/data-science-and-technology/idf/research/ceds-upmu-cyber-security/			

2016 - 2017	**DOE OE's Microgrid R&D**	**Micro-PMUs for advanced sensors integration and data analytics at the Navy Yard.**	LBNL and PIDC at the Philadelphia Navy Yard.
Objective: Demonstrate advanced analytics and algorithms utilizing micro-PMUs within a true microgrid environment.			
Link: https://gig.lbl.gov/news/grid-integration-group-receives-funding-doe			
2016 - 2018	**DOE ARPA-E**	**Plus-Up Project (continuation of above "Micro-Synchrophasors for Distribution Systems")**	CIEE/BECI, UC Berkeley, LBNL and LLNL, Doosan GridTech, Smarter Grid Solutions, PingThings.
Objective: Further develop and commercialize early diagnostic applications.			
2016 - 2018	**PCARI**	**Resilient Electricity Grids**	UC Berkeley and University of the Philippines, Diliman.
Objective: Deploy distribution PMUs with Dagupan Electric Cooperative in the Philippines, and demonstrate that the micro-PMU system is a cost-effective approach.			
Link: http://pcariofficial.blogspot.com/p/about-us_15.html			
2016	**DOE GMLC**	**GMLC 1.2.5 – Sensing & Measurement Strategy.**	LBNL, LLNL, and others.
Objective: Develop a cohesive strategy to develop and deploy sensing & measurement technologies for developing grid modernization metrics. Leverages micro-PMU advances.			
Link: https://energy.gov/under-secretary-science-and-energy/doe-grid-modernization-laboratory-consortium-gmlc-awards			
2016	**DOE GMLC**	**GMLC 1.4.9 - Integrated Multi Scale Data Analytics and Machine Learning.**	LBNL, LLNL, LANL, NREL, SNL, ANL
Objective: Multi-scale data analytics for distribution, leverage micro-PMU advances.			
Link: https://energy.gov/under-secretary-science-and-energy/doe-grid-modernization-laboratory-consortium-gmlc-awards			
2016	**DOE GMLC**	**Project 4: Project 4: Advanced Machine Learning for Synchrophasor Technology.**	LANL, BPA, JSIS, OPE Energy Corporation, Riverside Public Utilities
Objective: Develop a suite of new Grid-Modeling aware Machine Learning tools to monitor the transmission grid during normal operation and localize frequency events. Use aggregated micro-PMU measurements.			
Link: https://energy.gov/under-secretary-science-and-energy/doe-grid-modernization-laboratory-consortium-gmlc-awards			
2017 - 2018	**SENER and CONACyT**	**Bi-national Laboratory for the Intelligent Management of Energy Sustainability and Technology Education.**	UC Berkeley and ITESM, Universidad Tecnologico de Monterrey.
Objective: Deploy distribution PMUs with Comision Federal de Electricidad (CFE) in Mexico City.			

<table>
<tr><td>2017 - 2019</td><td>DOE ENERGISE</td><td>Phasor-Based Control for Scalable Solar Photovoltaic Integration</td><td>CIEE, UC Berkeley, Doosan GridTech, OPAL-RT, GridBright, U Michigan.</td></tr>
<tr><td colspan="4">Objective: Develop a new layered control framework for managing extremely high penetrations of solar generation and other Variable Energy Resources.</td></tr>
<tr><td colspan="4">Link: https://energy.gov/eere/sunshot/funding-opportunity-announcement-enabling-extreme-real-time-grid-integration-solar</td></tr>
<tr><td>2017</td><td>NASPI</td><td>Technical Report (NASPI-2017-TR-001) "Time Synchronization in the Electric Power System".</td><td>NASPI Time Synchronization Task Force.</td></tr>
<tr><td colspan="4">Objective: Identify and articulate what power system engineers and operators need to know about the role and emerging importance of high-quality timing sources in routine and mission-critical grid applications. Discuss micro-PMU grid application timing requirements.</td></tr>
<tr><td colspan="4">Link: https://www.naspi.org/</td></tr>
<tr><td colspan="4">INTERNATIONAL</td></tr>
<tr><td>2013 - 2017</td><td>Swiss National Science Foundation</td><td>Nano-Tera SmartGrid</td><td>EPFL-DESL & LCA2</td></tr>
<tr><td colspan="4">Objective: Produce and make publicly available a real-world micro-PMU data set for research and application development. First live demonstration of real-time (low latency, 20 ms refresh rate) state-estimation in a real-scale MV distribution grid, using advanced PMUs, PDC and state estimation processes developed by EPFL researchers.</td></tr>
<tr><td colspan="4">Links:
• https://smartgrid.epfl.ch
• https://infoscience.epfl.ch/record/203775</td></tr>
<tr><td>2012 - 2016</td><td>FP7-ICT (EU)</td><td>C-DAX: Cyber-secure Data and Control Cloud for Power Grids</td><td>Alcatel Lucent, iMinds, EPFL, EKUT, Radboud University Nijmegen, Liander, National Instruments, UCL, University of Surrey.</td></tr>
<tr><td colspan="4">Objective: C-DAX is a Cyber-secure DAta and Control Cloud for future power distribution networks based on an information-centric networking (ICN) architecture. C-DAX was validated in a distribution grid of Alliander (Dutch DSO), to demonstrate that the use of PMUs in combination with C-DAX and a public LTE network can help operators to better manage their grids by providing real-time situational awareness in a cost-efficient way.</td></tr>
<tr><td colspan="4">Links:
• http://www.cordis.europa.eu/project/rcn/106390_en.html
• https://infoscience.epfl.ch/record/222877</td></tr>
<tr><td>2014 - 2017</td><td>Swiss Federal Office of Energy (SFOE)</td><td>Real-time state estimation of the Lausanne 125 kV sub-transmission network using PMUs</td><td>EPFL-DESL, Services industriels de Lausanne (SiL), National Instruments</td></tr>
<tr><td colspan="4">Objective: validation of PMU-based real-time state estimation in the 125 kV sub-transmission grid of the city of Lausanne and demonstration of its capability to support hard real-time applications, such as power-system protections. Validation of a PMU-based fault location technique.</td></tr>
</table>

References

[1] FNET/GridEye web display. Operated by the Power Information Technology Laboratory at the University of Tennessee, Knoxville. Available: fnetpublic.utk.edu.

[2] Texas synchrophasor network. Baylor University. Available: https://web.ecs.baylor.edu/faculty/grady/Texas Synchrophasor Network.html

[3] A. Phadke and J. Thorp (2008). Synchronized Phasor Measurements and their Applications. New York: Springer.

[4] J. D. L. Ree, V. Centeno, J. S. Thorp, and A. G. Phadke, “Synchronized phasor measurement applications in power systems,” IEEE Transactions on Smart Grid, vol. 1, no. 1, pp. 20–27, June 2010.

[5] Y. Liu, Y. Liu, Y. Zhang, J. Guo, and D. Zhou, Wide Area Monitoring through Synchrophasor Measurement. John Wiley & Sons, Ltd, 2016. Available: http://dx.doi.org/10.1002/9781118755471.sgd029.

[6] A. Silverstein, M. Weimar, and J. Petersen, The value proposition for synchrophasor technology: Itemizing and calculating the benefits from synchrophasor technology use. North American Synchrophasor Initiative (NASPI) Report, October 2015.

[7] A. von Meier, E. Stewart, A. McEachern, M. Andersen and L. Mehrmanesh, “Precision Micro-Synchrophasors for Distribution Systems: A Summary of Applications.” IEEE Transactions on Smart Grid, July 2017. DOI 10.1109/TSG.2017.2720543.

[8] M. Jamei, E. Stewart, S. Peisert, A. Scaglione, C. McParland, C. Roberts, and A. McEachern (2016). Micro Synchrophasor-Based Intrusion Detection in Automated Distribution Systems: Toward Critical Infrastructure Security. IEEE Internet Computing, 20(5), 18-27.

[9] E. C. Kara, C. M. Roberts, M. Tabone, L. Alvarez, D.S. Callaway, and E. Stewart (2017). "Disaggregating Solar Generation From Feeder-Level Measurements". Submitted to Sustainable Energy, Grids and Networks.

[10] S. Peisert, C. McParland, J. Boverhof, R. Gentz, A. Elbashandy, and D. Gunter (May 31, 2017). LBNL Open Power Data, DOI 10.21990/C21599.

[11] G. Cavraro, R. Arghandeh, G. Barchi, and A. von Meier, "Distribution network topology detection with time-series measurements," in Innovative Smart Grid Technologies Conference (ISGT), 2015 IEEE Power & Energy Society, 2015, pp. 1-5.

[12] G. Cavraro, R. Arghandeh, K. Poolla, and A. von Meier, "Data-driven approach for distribution network topology detection," in Power & Energy Society General Meeting, 2015 IEEE, 2015, pp. 1-5.

[13] Saurav Talukdar, Deepjyoti Deka, Donatello Materassi, Murti Salapaka, "Exact topology reconstruction of radial dynamical systems with applications to distribution system of the power grid", 2017 American Control Conference (ACC) .

[14] R. Arghandeh, M. Gahr, A. von Meier, G. Cavraro, M. Ruh, and G. Andersson, "Topology detection in microgrids with micro-synchrophasors," in Power & Energy Society General Meeting, 2015 IEEE, 2015, pp. 1-5.

[15] D. Deka, M. Chertkov, and S. Backhaus, “Structure learning in power distribution networks,” in IEEE Transactions on Control of Network Systems, Feb 2017.

[16] William O’Brien, Eric Udren, Kamal Garg, Dennis Haes, and Bala Sridharan, “Catching Falling Conductors in Midair – Detecting and Tripping Broken Distribution Circuit Conductors at Protection Speeds”, 69th Annual Conference for Protective Relay Engineers, 2016, DOI https://doi.org/10.1109/CPRE.2016.7914881, Available at https://cdn.selinc.com/assets/Literature/Publications/Technical%20Papers/66_CatchingFalling_KG_20160215_Web3.pdf?v=20170706-140455.

[17] C.M. Roberts, C.M. Shand, K.W. Brady, E.M. Stewart, A.W. McMorran, and G.A. Taylor, Improving Distribution Network Model Accuracy using Impedance Estimation from Micro-Synchrophasor Data. http://bura.brunel.ac.uk/bitstream/2438/14118/1/Fulltext.pdf

[18] Pacific Gas and Electric information sheet on phase unbalance http://www.pge.com/includes/docs/pdfs/mybusiness/customerservice/energystatus/powerquality/voltage_unbalance_rev2.pdf (accessed December 2016).

[19] V. Arya, D. Seetharam, S. Kalyanaraman, K. Dontas, C. Pavlovski, S. Hoy, JR. Kalagnanam, “Phase identification in smart grids,” in IEEE Smart Grid Communications (SmartGridComm), 2011 IEEE, pp. 25-30.

[20] E.M. Stewart, C. Roberts, A. Liao, A. von Meier, O. Ardakanian and K. Brady, and A. McEachern, “Predictive distribution component health monitoring with distribution phasor measurement Units.”

[21] KX Perez, WJ Cole, JD Rhodes, A Ondeck, M Webber, M Baldea, TF Edgar, “Nonintrusive disaggregation of residential air-conditioning loads from sub-hourly smart meter data,” Energy and Buildings. 2014 Oct 31, 81:316-25.

[22] CL. Holcomb, “Disaggregation of Residential Electric Loads Using Smart Metered Data,” Masters Thesis, University of Texas at Austin, 2011.

[23] Mayhorn ET, Sullivan GP, Butner RS, Hao H, Baechler MC., “Characteristics and performance of existing load disaggregation technologies,” Pacific Northwest National Laboratory (PNNL), Richland, WA (United States), 2015 Apr 10.

[24] Zoha A, Gluhak A, Imran MA, Rajasegarar S., Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. Sensors. 2012 Dec 6, 12(12):16838-66.

[25] Kleiminger W, Beckel C, Staake T, Santini S. Occupancy detection from electricity consumption data. Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings 2013 Nov 11 (pp. 1-8). ACM.

[26] McLaughlin S, McDaniel P, Aiello W. Protecting consumer privacy from electric load monitoring. InProceedings of the 18th ACM conference on Computer and communications security 2011 Oct 17 (pp. 87-98). ACM.

[27] Y. Zhou, R. Arghandeh, I. Konstantakopoulos, S. Abdullah, A. von Meier, and C. Spanos, "Abnormal event detection with high resolution micro-PMU data." In 2016 Power Systems Computation Conference (PSCC), Genoa, 2016.

[28] J. Lee, "Automatic fault location on distribution networks using synchronized voltage phasor measurement units." In Proceedings of the ASME 2014 Power Conference, POWER2014-32231, Baltimore, MD, July 2014.

[29] S. Jamali and A. Bahmanyar, "A new fault location method for distribution networks using sparse measurements." Electrical Power and Energy Systems 81, 459-468 (2016).

[30] C. Orozco-Henao, A. Bretas, R. Chouhy-Leborgne, A. Herrera-Orozco, and J. Marín-Quintero, "Active distribution network fault location methodology: A minimum fault reactance and Fibonacci search approach." Electrical Power and Energy Systems 84, 232-241 (2017).

[31] L. Schenato et al., "Bayesian linear state estimation using smart meters and PMU measurements in distribution grids," in Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm), Venice, Italy, Nov 2014, pp. 572-577.

[32] E. Cotilla-Sanchez, P. D. Hines, and C. M. Danforth, "Predicting critical transitions from time series synchrophasor data," IEEE Transactions on Smart Grid, vol. 3, no. 4, pp. 1832–1840, 2012.

[33] M. Chertkov, S. Backhaus, K. Turitsyn, V. Chernyak, and V. Lebedev, "Voltage collapse and node approach to power flows: analysis of a feeder line with static disorder in consumption/production," arXiv preprint arXiv:1106.5003, 2011.

[34] D. Podolsky and K. Turitsyn, "Random load fluctuations and collapse probability of a power system operating near codimension 1 saddle-node bifurcation," in Power and Energy Society General Meeting (PES), 2013 IEEE. IEEE, 2013, pp. 1–5.

[35] D Deka, A Zare, A Lokhov, M Jovanovic, and M Chertkov, "State and Noise Covariance Estimation in Power Grids using Limited Nodal PMUs," IEEE GlobalSIP (Smart Grid Infrastructure), 2017.

[36] A. von Meier, D. Culler, A. McEachern and R. Arghandeh, "Micro-synchrophasors for distribution systems," ISGT 2014, Washington, DC, 2014, pp. 1-5. doi: 10.1109/ISGT.2014.6816509.

[37] Y. Seyedi, H. Karimi, and J. Guerrero, "Centralized disturbance detection in smart microgrids with noisy and intermittent synchrophasor data," IEEE Transactions on Smart Grid, Early Access, DOI: 10.1109/TSG.2016.2539947, 2016.

[38] L. Zhan, J. Zhao, J. Culliss, Y. Liu, Y. Liu and S. Gao, "Universal Grid Analyzer design and development," 2015 IEEE Power & Energy Society General Meeting, Denver, CO, 2015, pp. 1-5.

[39] L. Zhan, Y. Liu, J. Culliss, J. Zhao and Y. Liu, "Dynamic Single-Phase Synchronized Phase and Frequency Estimation at the Distribution Level," in IEEE Transactions on Smart Grid, vol. 6, no. 4, pp. 2013-2022, July 2015.

[40] M. Wu and L. Xie, "Online detection of low-quality synchrophasor measurements: A data driven approach," IEEE Transactions on Power Systems, Early Access, DOI: 10.1109/TPWRS.2016.2633462, 2016.

[41] IEEE Team. "IEEE Standard Requirements for Instrument Transformers," IEEE Std C57.13-2016 (Revision of IEEE Std C57.13-2008), vol., no., pp.1-96, June 29 2016. doi: 10.1109/IEEESTD.2016.7501435, URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7501435&isnumber=7501434.

[42] B. Gou, "Optimal placement of PMUs by integer linear programming," IEEE Trans. Power Syst., vol. 23, no. 3, pp. 1525–1526, Aug. 2008.

[43] B. Gou and A. Abur, "An improved measurement placement algorithm for network observability," IEEE Trans. Power Syst., vol. 16, no. 4, pp. 819–824, Nov. 2001.

[44] R. Emami and A. Abur, "Robust measurement design by placing synchronized phasor measurements on network branches," IEEE Trans. Power Syst., vol. 25, no. 1, pp. 38–43, Feb. 2010.

[45] T. Baldwin, L. Mili, M. Boisen, Jr, and R. Adapa, "Power system observability with minimal phasor measurement placement," IEEE Trans. Power Syst., vol. 8, no. 2, pp. 707–715, May 1993.

[46] F. Magnago and A. Abur, "A unified approach to robust meter placement against loss of measurements and branch outages," IEEE Trans. Power Syst., vol. 15, no. 3, pp. 945–949, Aug. 2000.

[47] F. Aminifar, C. Lucas, A. Khodaei, and M. Fotuhi-Firuzabad, "Optimal placement of phasor measurement units using immunity genetic algorithm," IEEE Trans. Power Del., vol. 24, no. 3, pp. 1014–1020, Jul. 2009.

[48] Li, X., Scaglione, A., & Chang, T. H. (2014). A framework for phasor measurement placement in hybrid state estimation via Gauss–Newton. IEEE Transactions on Power Systems, 29(2), 824-832.

[49] M. Jamei *et al*., "Anomaly Detection Using Optimally-Placed µPMU Sensors in Distribution Grids," in *IEEE Transactions on Power Systems*, vol. PP, no. 99, pp. 1-1. doi: 10.1109/TPWRS.2017.2764882

[50] M. P. Andersen, S. Kumar, C. Brooks, A. von Meier, and D. E. Culler, "DISTIL: Design and implementation of a scalable synchrophasor data processing system," in Proc. IEEE Conf. Smart Grid Commun., Miami, FL, USA, 2015, pp. 271–277.

[51] M. P. Andersen and D. E. Culler, "BTrDB: Optimizing storage system design for timeseries processing," in Proc. 14th USENIX Conf. File Storage Technol. (FAST), Santa Clara, CA, USA, 2016, pp. 39–52.

[52] NASPI Time Synchronization Task Force Technical Report "Time Synchronization in the Electric Power System" (NASPI-2017-TR-001), March 30th, 2017.

[53] R. S. Singh, H. Hooshyar, L. Vanfretti, "Laboratory test set-up for the assessment of PMU time synchronization requirements," in Proc. IEEE PowerTech Conference, Eindhoven, Netherlands, June 29-July 2, 2015. DOI: 10.1109/PTC.2015.7232731. URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7232731&isnumber=7232233

[54] ALSETLab Github Repository, "IRIG-B Signal Generator Code in Matlab/Simulink for use with Real-Time Simulators", 2017, https://github.com/ALSETLab/IRIG-B_for_RT.

[55] R. S. Singh, H. Hooshyar, L. Vanfretti, "Assessment of time synchronization requirements for phasor measurement units," in Proc. IEEE PowerTech Conference, Eindhoven, Netherlands, June 29-July 2, 2015.

[56] Clement, P., & Johnson, W. (1960). Electrical Engineering Science. New York, NY: McGraw-Hill.

[57] Scharf, L., & Behrens, R. (1990). A First Course in Electrical and Computer Engineering. Reading, MA: Addison-Wesley.

[58] Moura, L., & Darwazeh, I. (2005). Introduction to Linear Circuit Analysis and Modelling. Oxford, England: Newnes.

[59] "IEEE Standard For Synchrophasor Measurements For Power Systems," IEEE Standard C37.118.1-2011 (Revision of IEEE Std C37.118-2005), pp. 1–61, Dec 2011.

[60] Stevens, S. (1946). On the theory of scales of measurement. Science, 103, 677-680.

[61] Ellis, B. (1966). Basic concepts of measurement. London: Cambridge University Press.

[62] Luce, R., & Tukey, J. (1964). Simultaneous conjoint measurement: J Math Psych, 1, 1-27.

[63] IEC 60044-1, International Electrotechnical Commission (IEC). 2003. Instrument Transformers -- Part 1: Current Transformers.

[64] Kirkham, H. (2016). The Measurand: the problem of Frequency. Proceedings of the International Instrumentation and Measurements Conference, I&M Society. Taipei: IEEE.

[65] Kirkham, H. and Riepnicks, A. Measurement of Phasor-like Signals. PNNL-25643, June 2016. Available: https://www.naspi.org/sites/default/files/reference_documents/pnnl_25643.pdf