# DESIGN, IMPLEMENTATION, AND EVALUATION OF NAPALI: A NOVEL DISTRIBUTED SENSOR NETWORK FOR IMPROVED POWER QUALITY MONITORING.

## A THESIS PROPOSAL SUBMITTED TO MY COMMITTEE IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

## DOCTOR OF PHILOSOPHY

IN

## COMPUTER SCIENCE

DECEMBER 2018

By

Sergey Negrashov

Thesis Committee:

Philip Johnson, Chairperson Peter-Michael Seidel Edo Biagioni Lipyeow Lim Matthias Fripp

Keywords: Sensor Network, Distributed Sensing, Grid Computing, Power Quality

Copyright © 2018 by Sergey Negrashov

## ABSTRACT

Today's big data world is heavily relied on to bring precise, timely, and actionable intelligence, while being burdened by the ever increasing need for data cleaning and preprocessing. While in the case of ingesting large quantity of unstructured data this problem is unavoidable, when it comes to sensor networks built for a specific purpose, such as anomaly detection, some of that computation can be moved to the edge of the network. This thesis concerns the special case of sensor networks tailored for monitoring the power grid for anomalous behavior. These networks consist of meters connected to the grid across multiple geographically separated locations, while monitoring the power delivery infrastructure with the intent of finding deviations from the nominal steady state. These deviations, known as power quality anomalies, may originate, and be localized to the location of the sensor, or may affect a sizable portion of the power grid. The difficulty of evaluating the extent of a power quality anomaly stems directly from their short temporal and variable geographical impact. I propose a novel distributed power quality monitoring system called Napali which relies on extracted metrics from individual meters and their temporal locality in order to intelligently detect anomalies and extract raw data within temporal window and geographical areas of interest. The results of this research should be useful in other disciplines, such as general sensor network applications, IOT, and intrusion detection systems.

# TABLE OF CONTENTS

Abstract												
$\mathbf{Li}$	List of Tables											
Li	st of	Figures	vii									
1	Intr	$\mathbf{r}$ oduction $\ldots$	1									
	1.1	Overview of power grids	1									
	1.2	Edge computing approach to anomaly detection	2									
	1.3	Napali: hybrid edge computing for anomaly detection.	4									
	1.4	The problem of power quality	6									
	1.5	Edge computing approach to power quality monitoring.	8									
	1.6	Thesis claim and evaluation	10									
<b>2</b>	Rela	ated Work	14									
	2.1	Edge computing	14									
	2.2	Distributed Power Quality Monitoring	17									
	2.3	Anomaly detection in Power Quality Monitoring Networks.	18									
3	Ope	en Power Quality	20									
	3.1	OPQ Box	20									
	3.2	OPQ Makai	26									
		3.2.1 Triggering Broker	26									
		3.2.2 Acquisition Broker	27									
		3.2.3 Acquisition Service	28									
	3.3	OPQ Mauka	28									

	3.4	OPQ	View
4	Eva	luatior	$1 \dots \dots$
	4.1	OPQ I	Box Validation
	4.2	Napali	Validation
		4.2.1	Napali Bandwidth usage
		4.2.2	Sink processing requirement under the Napali Framework
		4.2.3	Effects of latency in the Napali framework
		4.2.4	Temporal locality triggering of the Napali framework
		4.2.5	Sub-threshold Data Acquisition
	4.3	Napali	Framework in Other Domains
	4.4	Schedu	ıle
Bi	bliog	graphy	

# LIST OF TABLES

# LIST OF FIGURES

1.1	Comparison of the three event detection methodology across three metrics. Methods are as follows: naive method (2), self triggering (1), Napali, hybrid solution (3) $\ldots$	4
1.2	RMS voltage waveform generated form the consumer side of the meter under various conditions. All waveform were recorded using the OPQ Box 2	7
1.3	Power quality anomaly propagation example	10
2.1	Projected number of IOT devices worldwide.[1]	14
2.2	Content delivery network architecture. As described in the Google pattent.[6] $\ldots$	15
2.3	Fog computing use in transportation. The bus cloudlet provides a cache for common data such as commuter schedules and traffic information, while routing other queries to the Internet.[25]	16
3.1	Block diagram of the OPQ Power Quality monitoring system	20
3.2	(a) OPQ Box2 block diagram and (b) production OPQ box ready for deployment	21
3.3	Block diagram of the OPQ Box 2 software stack.	22
3.4	Filters used for mains frequency calculation. (a) Downsampling filter gain. (b) Downsampling filter impulse response. (c) Lowpass filter gain. (d) Lowpass filter impulse response	23
3.5	Filters used for mains frequency calculation. (a) Downsampling filter gain. (b) Downsampling filter impulse response. (c) Lowpass filter gain. (d) Lowpass filter impulse response.	24
3.6	A lightning strike recorded by two OPQ Box 2 devices separated by 10 miles. (a) A lightning strike manifested as $V_{rms}$ dip which lated 11 cycles. (b) As a consequence of using NTP these devices have $\frac{1}{2}$ cycle mismatch in reported timestamps	25
3.7	Block diagram of the OPQ Makai.	27
3.8	Block diagram of the OPQ Mauka.	29

3.9	Screenshot of a recent OPQ View build	•	•	 •	•	•	•	•	•	•	 30
4.1	University of Hawaii at Manoa power delivery infrastructure.										 32

## CHAPTER 1 INTRODUCTION

Power quality research is a subset of power distribution research which focuses on studying deviations from nominal power grid operating conditions. Devices connected to the power grid, as well as the distribution equipment expect a certain frequency, voltage and harmonic content of the voltage waveform they operate on. While most equipment maintains some hysteresis with respect to deviations from the nominal, large enough deviations may cause equipment failure and instability in the power grid as a whole. In a practical sense, power quality monitoring concerns itself with monitoring, collecting and analyzing power quality anomalies on a live and functioning grid. In some cases, for example when performed by the utility, this information is used to make realtime decisions, to maintain the stability of the power grid. However, data collected by the power quality monitoring equipment can also be used to diagnose local power quality problems, or to further power quality research. Power quality monitoring fits very well into the paradigm of remote sensing and sensor networks, particularly into the newly emerging field of edge computing. Edge computing goes beyond the naive approach of transmitting the entirety of the collected data from the sensor location, and extends it by either feature extracting, preprocessing or filtering the data at the computing node itself. This research project will investigate the design, implementation, and evaluation of a novel edge computing architecture called Napali which combines feature extraction at the edge level and two way communication between the sink and the edge node. I will evaluate Napali in part by implementing it in the power quality monitoring domain.

## 1.1 Overview of power grids

Modern power grids are hierarchically structured. Higher voltage is useful for transporting electricity over long distances, connecting cities and towns to power generation facilities. Transmissions lines of 100kV and above are used to minimize losses in long distance runs, since the same amount of power can be transmitted using much lower current, and thus much more efficiently, then the comparable low voltage line. Close to the point of distribution, transmission voltage is stepped down to 1kV-40kV range using large power transformers. This is done because the losses incurred in the final leg of transmission are minimal, while extremely high voltage equipment is expensive and requires special precautions. [24] Finally, at the consumer level the voltage level is stepped down once more to the household voltage, for example  $120V_{ac}$  for North America. It is important to note that voltage across every part of the power grid is synchronized to a phase and frequency set by the utility. This allows multiple power producers to contribute to electricity generation without interfering with each other. [4] In North America the 60Hz utility frequency is used as the baseline, and its long term stability is guaranteed by the power company. How close the power AC frequency is to the nominal value is a measure of how closely the electricity demand is

balanced by the electricity generation.

Traditional power generation sources involves applying mechanical torque to an alternating current generator. If the load on the generator increases without increasing the torque, it will slow down the generator and thus the utility frequency decreases. Similarly, if the demand drops but the torque is not decreased, the frequency of generated power will increase. Even small deviations in frequency can have adverse effects on equipment which runs synchronous to the power grid, such as synchronous electric motors and other industrial equipment.[15] Nonlinear loads, or loads that don't draw a consistent amount of current through out an AC cycle, are highly prevalent in today's power grid. These devices contribute to the harmonic noise in power system in both current and voltage waveform. This effect, known as harmonic distortion, can have various unintended consequences on the power distribution system and connected devices. The current harmonic distortion affects the efficiency of the distribution network, while voltage harmonics may propagate across the power distribution infrastructure and affect neighboring devices.[16] Distributed renewable generation may also create untended harmonics. Distributed generators are commonly DC systems, which utilize inverters to generate in phase AC waveform to feed into the power grid. Depending on the inverter design the AC waveform may have spurious harmonics present.[15]

Large and sudden changes in load to generation ratio change do not immediately impact the utility frequency due to the rotational inertia of the large generation systems. Instead it will cause the line voltage to change proportional to the load until the generation can catch up. If the load suddenly increases, caused for example by a large motor stall, grid voltage will experience a sharp drop, known as a sag. Similarly a large load drop will cause an voltage increase, called a swell. Voltage sags and swells propagate throughout the entire grid infrastructure, however the dynamics of the power grid are quite complex, and hard to predict. For example a voltage sag on one sub-transmission chain may manifest and as a voltage swell in another. Finally very fast changes in load, such as short circuits, opening and closing of re-closers, and lightning strikes manifest as voltage transients. Voltage transients are energetic short-lived swells on the order of a single AC cycle, which can travel across the distribution grid. Transients may interfere with sensitive grid connected equipment, as well as trigger protection equipment such as uninterpretable power supplies, and other over-voltage protection devices. Transients, harmonic distortion, and RMS fluctuations and their combinations make up the majority of power quality problems which affect the voltage waveform in the grid connected devices. [3] All of these issues can cause power quality problems, as will be discussed further in Section 1.4.

## **1.2** Edge computing approach to anomaly detection

Edge computing is an emergent field in distributed systems. Edge computing is a consequence of ever decreasing power consumption of computational devices found on the sensor nodes, as well as incremental improvements in battery technology. With ever-increasing computational capabilities in sensor networks, it becomes possible to process and store the acquired data on the device itself, as opposed to the centralized sink. Thus the idea of edge computing leverages available computing power of the sensor node to allow for smarter distributed sensing. Edge computing with respect to remote sensing allows for several new approaches to anomaly detection.

Anomaly detection is a common topic across many disciplines and domains. In cyber-security research, anomalous network traffic and program behavior is often indicative of malicious behavior. In seismic monitoring, anomalies in ground vibrations may be precursors to an earthquake or a volcano eruption. In observational astronomy, anomaly detection is used for detection of transient events such as gamma ray bursts. Sensor networks are commonly tasked with anomaly detection and must often act on them. Traditionally, stringent constraints on power consumption of battery powered wireless sensor network nodes mean that low bandwidth and low complexity methods are preferred. Furthermore, many sensor networks are often hindered by local noise, thus requiring higher level filtering in order and in network processing to determine if an anomaly has occurred. If the signal to noise of the local measurements is quite high this problem becomes trivial: one simply collects all the distributed measurements if one or more of the measurements indicates an anomaly. Unfortunately, in the real world such problems are rare and instead the distributed signal is dominated by extraneous local noise. For example individual seismic sensors can't distinguish between a global anomaly such as an earthquake and local noise such as vibration caused by a passing semi-truck.

The problem of global anomaly detection with distributed sensing has been thoroughly explored in self organizing wireless sensor networks. However, these approaches are insufficient in the paradigm of edge computing. Edge computing relies on Internet for transport, and thus the cost of communicating with the local sink and the local node is similar. Indeed in some cases it is impossible to achieve node-to-node communication without an intermediary due to firewalls, and other security mechanisms. In my research I will only consider approaches which rely on a sink node to facilitate anomaly detection.

There are several solutions for dealing with the local noise problem in an Internet-enabled sensor network. A naive solution is to simply record every distributed measurement to a centralized data sink. This sink, as well as the infrastructure down stream of it has a view of the entire state of the system and can thus detect anomalies using either real-time or batch processing. An alternative to sending all of the distributed measurements to the sink is to let the individual sensors decide which temporal regions of the measurement constitute an anomaly. This approach, called self triggering, has a benefit of the reducing the bandwidth constraints for each sensor without the requirement for two way communication between the sensor and the sink.

## **1.3** Napali: hybrid edge computing for anomaly detection.

I propose a framework called Napali which combines the strengths of the previously mentioned methods. In Napali, each sensor node maintains a two way communication channel with the sink, as well as a temporal window containing all the recent data each device has collected. Each sensor's on board processing is used to extract features from the collected measurements, and only these features, instead of the entire measurement set are forwarded to the sink for processing. The sink has a low resolution view of the state of the entire sensor network. While this may not be enough for rigorous anomaly analysis, properly feature extracted data from every node should be enough to detect the occurrence of an anomaly. Finally if an anomaly is detected the sink will request high resolution data from all of the affected devices and their neighbors.



Figure 1.1: Comparison of the three event detection methodology across three metrics. Methods are as follows: naive method (2), self triggering (1), Napali, hybrid solution (3)

The naive method provides the best detection ability and the smallest node computation requirement. However, it does so at the cost of the largest bandwidth consumption. The self triggering method has the lowest bandwidth consumption of the three. The disadvantage of this method is twofold. First of all in order to maintain a high detection rate a reasonably low threshold for anomalies has to be used. This may cause a large number of false positives due to local noise and sensor noise. Secondly, global anomalies will often diminish in magnitude as a function of distance from the epicenter, thus far removed sensor nodes may completely miss events. These low threshold events may be invaluable for reconstructing the dynamic of the anomaly propagation, however they will be missed by the detection algorithms.

Napali's hybrid approach provides the opportunity for much better anomaly detection rate and background noise rejection, by correlating the features that are computed on the sensor nodes. Napali has moderate bandwidth consumption, however the bandwidth consumption is tunable since the features can be computed for varied temporal windows length of which can be adjusted in real time. Napali does requires the highest sensor node computing power, since not only does it need to extract the triggering features from from raw data, it needs to maintain a buffer of sensor measurements to send to the sink if an anomaly is detected. The downside of the Napali framework, is the requirement for two way communication between the sink and the sensor node. In order to participate in event detection Napali sensors need to be able to both send raw and feature extracted data, as well as receive the control signals from the sink. In contrast, the naive and self triggered approaches only require a one way link to send the raw data to the sink. As such, standalone devices can be retrofired to participate in distributed event detection in the naive and self triggered approaches, while Napali requires custom devices specifically tailored for cooperative event detection.

I claim that there are several benefits to the Napali method compared to the naive and self triggered approaches.

- 1. Bandwidth usage is minimized: Instead of sending the entirety of raw data, only extracted features are sent. This features will have a tiny fraction of the bandwidth requirement when compared to raw waveforms. Furthermore, the temporal window which encompasses a single feature can be adjusted in real time. Thus as soon as an anomalous behavior is observed in a subset of sensors, this window can be adjusted for a finer grained feature extraction.
- 2. Effects of latency are minimized: Even at 1Msample/second at 16bits of resolution, the memory requirement to store 5 minutes of raw waveform without compression are on the order of 512MB, which is well within the realm of cheap single board computers. With compression specifically suited to the signal of interest, the memory requirement can be reduced even further. This gives the triggering stream sink plenty of time to respond to the anomalies in the data and request raw waveform from the monitoring devices.
- 3. Sink processing requirements are minimized: Since most of the feature extraction is already performed at the device level the triggering stream sink computational resources can be minimized. With the advent of IOT, the computational capacity of the edge devices is quickly growing. Napali exploits these resources, and thus minimizes the sink computational requirements.
- 4. **Sub-threshold data acquisition:** The triggering stream sink makes the decision to acquire raw waveform from sensor nodes. This allows researchers to collect data from devices which were only mildly affected or not affected by the disturbance. This provides new possibilities for investigation of disturbance propagation across the sensed area.
- 5. Power failure resiliency: In the case of the complete power failure or communication blackout, if the monitoring device has a battery backup capability, each sensor has a record of the entire raw waveform leading up to the power interruption. Prior to shutdown the sensor node will transfer all of the raw data from the volatile memory to on-board permanent storage. Once the power or communication channel is restored, select portions of the buffer may be sent back to the data sink. This creates an additional layer of resiliency for the anomaly detection network.

- 6. Flexible Privacy: Anomalies which were only observed at a single point are most likely local noise and pose little value for global state monitoring. It's up to the user to decide how to process disturbances which affect their sensor. For example user may choose to record a full waveform, only certain features, or record nothing at all. Secondly, if the saturation of the device is high enough only a subset of them would need to send a triggering stream for event detection, while the rest will be used for acquiring raw waveform for small temporal regions containing global events.
- 7. **Temporal Locality:** By exploiting the temporal locality it is possible to ascertain the geographical extent of the anomaly with only coarse features. This allows for a simple robust algorithm which may be deployed at the sink node for anomaly detection.

## 1.4 The problem of power quality

Power quality monitoring along with other smart grid domains is a field well suited for distributed sensor network monitoring. [14] [19] As the power grid moves from centralized generation with a few centralized power plants to distributed generation with residential power production, a distributed consumer level monitoring system becomes ever more important. Traditionally utilities do not monitor the quality of power besides the consumption beyond the substation level. This is due to the fact that the last opportunity that the utility has to regulate and filter the grid voltage in the hierarchical power distribution is at the substation, or neighborhood level. This methodology is unsustainable, as residential renewable energy production if not properly monitored and controlled will have an adverse effect on the overall grid stability. Furthermore, lack of residential monitoring may lead to dangerous conditions such as islanding, where an otherwise powered down grid may have a small powered island sustained only by the unmonitored residential renewable generation. Finally, residential power quality monitoring gives utility costumers an opportunity to evaluate the quality of power delivered to their household. As consumer electronics are becoming more and more complex, they become more sensitive to the power anomalies. Grid monitoring is traditionally the responsibility of the utility, however in most cases utilities only have to disclose power interruptions lasting several minutes. Small interruptions, and partial interruptions such as voltage sags, swells, frequency fluctuations and transients will often go undisclosed by the utility and unnoticed by the consumer, but may cause premature failure in grid connected electronic devices. It is in the best interest of the consumers to monitor the quality of the power that is delivered to them, meanwhile the same same monitoring system will also allows researchers and utilities to monitor the entirety of smart grid.

Residential power quality monitoring presents a number of issues. First of all, it is difficult to discern which side of the utility meter a power disturbance came from. Consider a voltage sag generated by a 1kW hot plate shown in Figure 1.2a. Since the output impedance of the power grid is



(c) Line voltage recorded over 24 hours in a residential household with photovoltaic installation.

Figure 1.2: RMS voltage waveform generated form the consumer side of the meter under various conditions. All waveform were recorded using the OPQ Box 2.

non-zero a high powered device can cause a significant voltage sag affecting every device connected to the same circuit. Secondly, recoding the voltage waveform resulting from non linear load can result in privacy leaks for the end user. Consider the voltage waveform produced by a PC running a video game as shown in Figure 1.2b. Throughout the game loading process the power load varies based on which components are in use. Furthermore, regions with CPU load, harddisk load and video card load can be readily identified by measuring the resulting voltage sag. Recording these event's has an adverse effect on end-users privacy and offers no immediate benefit in studying grid stability. Finally distributed power generation such as rooftop solar has significant effect on the residential voltage waveform. Consider the voltage waveform shown in Figure 1.2c. This waveform was recorded over 24 hours in a household with a rooftop photovoltaic installation. Similar to the voltage sag case since the power grid impedance is non-zero residential power generation will cause a voltage swell during peak hours of sunlight as evident by the waveform. Combined with the global voltage sag during hours of peak demand (3pm) combined with the lack of PV production during that time, photovoltaics installations can result in a daily  $10V_{rms}$  swing. Residential power quality monitoring can be accomplished via a distributed sensor network made up of power quality monitoring devices with high degree of penetration across the end points of a power grid. Furthermore, in order to monitor the dynamics of the entire power grid via residential line voltage monitoring it is imperative to monitor across multiple locations simultaneously. This combined with temporal and spatial correlations of data produced by the sensors allows for identification of grid wide anomalies without a high rate of false positives. Additionally, not all power quality events affect the entire grid, due to the hierarchical structure of the power distribution. The final requirement for grid-wide monitoring is a novel distributed event detection system.

## 1.5 Edge computing approach to power quality monitoring.

IEEE1159 standard describes the techniques for single location power quality monitoring. For transient monitoring it suggests a sampling rate of at least 7680 samples/second, up to 1 Megasample/second. This implies that if a power quality event detection system relies on raw data from all monitors it would do so at a very large bandwidth cost. At 20 Ksamples/second at 16bit per samples a single monitoring device would generate 300Kb/second. Several thousand devices could easily saturate at 1Gb link with no obvious benefit. Collecting and recording all of the raw waveform data from residential power quality meters for batch processing presents some privacy concerns as outlined above. On board event detection methodology allows the measurement devices to select which temporal regions of measurements constitute an event. This is a perfect strategy from the consumer's perspective, since it would allow for recording of power quality information which directly impact their residence. As mentioned in Section 1.2, this method relies on a threshold based approach where every device has a compute several metrics from the raw waveform and compares then to preprogrammed threshold values. Metrics such as  $V_{rms}$ , fundamental frequency and THD are easily adapted for single point power quality measurements. Temporal regions during which these metrics surpass a threshold are considered by the device as a power quality event, and thus are recorded.

The problem with the method outlined above is that grid-wide power quality events do not affect the entire grid in the same way. For example due to the grids hierarchical structure a voltage sag on one sub-circuit can manifest as as a sag of a different magnitude or even a swell on another.[9] This may result in a situation where some of the monitoring devices will not consider a power quality anomaly as an event, because it did not surpass the metric threshold, and simply ignore it. From the research perspective, however, it is important to get raw data from all of the affected devices not just the ones that were the most affected. This makes a hybrid centralized/decentralized event acquisition strategy more attractive for distributed residential power quality monitoring. In this scheme all monitoring devices use local processing resources to feature extract the incoming waveforms while storing them locally for several minutes. A sink collects all of the extracted features and looks for anomalies which are present in the feature data stream which we will refer to as triggering stream. If an anomaly is present in only a single device, it is highly probable that the disturbance occurred in the residence. Depending on the user's privacy preference, raw data for a single device anomaly can be be recorded for later analysis, or in case of a highly privacy conscious user, ignored. On the other hand if the triggering stream shows an anomaly temporally collocated across multiple devices, the entire network or a subset of the network may be queried for raw waveform data for a temporal region which corresponds to the disturbance in the triggering stream.

The main disadvantage of this method is that while there are plenty of power quality event detection methodologies for single location, there has been little development in the distributed event detection methods and metrics. Two problems are quite similar, indeed one may use the same metrics for distributed event detection as with the single point power quality monitoring. However, it's also important to consider temporal locality of anomalies detected across multiple devices and effects of device synchronization. Power quality anomalies such as voltage sags and transients will propagate through the transmission lines at the speed of light, however due to the non-linear elements which make up the power-line junctions, a certain temporal spread in event detection across multiple locations is expected. For large power grids such continental United States grids, large frequency fluctuations propagate in a highly nonlinear ways. In these cases the event propagation is limited by the inherent rotational inertia of the power generation systems, and the speed at which the grid protection elements such as reclosers and circuit breakers operate. Regardless, the closer local anomalies are detected in time, the more likely are they to be a result of a gridwide event. Unfortunately, it is unfeasible to perfectly temporally synchronize the distributed power quality monitors. While methods such as GPS can in principle provide synchronization of up to 10ns jitter across a large geographical region, they require a line of sight to the sky, and add a non-trivial cost to the bill of materials for every power quality meter. Furthermore, GPS is prone to losing signal depending on atmospheric conditions, and can be very sensitive to fluctuations in the power supply voltage, a critical time in power anomaly detection. An alternative to GPS is Network Time Protocol. Network time protocol can provide timing synchronization on the order of 10ms across Internet, which is on the order of  $\frac{1}{2}$  of a grid cycle. NTP performance could be further improved by using geographically close time servers which are themselves synchronized via GPS. Consider a situation where two devices are located in household which experience a local 100mspower quality disturbance every 10 minutes. Even with a 10ms synchronization jitter, it will take on average 21.5 days before the two disturbances are observed within 20ms of each other. If a third device is introduced, it is highly unlikely that all 3 would observe unrelated local anomalies within 20ms of each other over the lifetime of the power quality monitoring network. This implies that combination of temporal and threshold based correlation on the feature extraction data would allow one to build a robust residential based power grid monitoring system which would yield a very low rate of false positives.



Figure 1.3: Power quality anomaly propagation example.

## **1.6** Thesis claim and evaluation

Today's big data world is plagued with the issues of data cleaning and validation, even though it's being relied on for timely, accurate and actionable intelligence. With large ingress of unstructured data these issues are unavoidable, and preprocessing will remain a large portion of the analysis workload. However, in the case of sensor networks designed for a specific purpose, the tasks of anomaly detection can be pushed to the edge of the network using the Napali methodology. The claim of this thesis is that Napali provides novel architecture that is both a feasible solution to the problem of distributed power quality monitoring and which also provides significant benefits over the two standard alternative architectures (all computation/storage at nodes, all computation/storage at the sink).

To evaluate the feasibility of the Napali framework, I propose to implement it as part of the Open Power Quality(OPQ) system and apply it to the problems of power quality monitoring. Combined with higher level anomaly analysis, Napali provides important services for a Open Power Quality power quality monitoring network. This network is made up of a group of monitoring devices as well as a centralized data sink server. This system will be deployed for testing at the University of Hawaii at Manoa campus, by deploying power quality monitors across at least 20 buildings on campus. The University of Hawaii at Manoa campus is a unique testbed for such a network, since entire campus is isolated to its own microgrid, connected to the municipal Oahu grid via a 46kV feeder. Furthermore, the University of Hawaii has deployed a set of smart power monitors at the key positions in the grid, which can be used as a state of the art ground truth for evaluation of OPQ performance.

The system relies on a custom residential power quality monitor called OPQ Box, designed specifically for distributed monitoring using the Napali framework. Instead of performing local analysis on the voltage waveform with the aim of PQ anomaly detection, or forwarding all the recording measurements to the centralized sink, OPQ boxes computes a small subset of features on the input voltage waveform. These features are then forwarded to the Napali framework's centralized sink which performs the anomaly detection on reduced data, while the raw waveform is retained for a short time on the OPQ Box. If the sink determines that a possible anomaly has occurred, a request is sent to the affected and nearby devices for raw data.

The goal of Napali is not to provide a low rate of false positives for a particular type of a power quality disturbance. Indeed, once the raw data is acquired by the sink, filtering through the potential anomalies is trivial using well established methods. Instead, the focus of my detection system is balancing low bandwidth required for detection with the low rate of false negatives. Furthermore, monitoring at the leaf edges relies on the hierarchical nature of the power grid in order to ascertain the state of the entire power generation and delivery system. As noted in the literature, power quality disturbances tend to propagate down the hierarchy as shown in Figure 1.3.

Consider a lightning strike on a hypothetical 12kV feeder line in a hierarchical power grid. The directly affected devices will be the ones downstream from the disturbance. These devices will experience the most severe effects, most notably transients, as they propagate throughput the power delivery infrastructure. The indirectly affected devices will expedience a power anomaly mainly attributed to the power production entities trying to compensate for the large disturbance caused by the lightning strike. Thus, monitoring of the leaf edges of the power delivery system can in principle provide insights into the disturbances that originate deep inside the power distribution network.

The Open Power Quality system is designed to be a test bed for development of new power quality detection and analysis algorithms. Its introduction will facilitate development of new techniques and methods for studying power system, by utilizing the Napali framework as the main anomaly detection methodology. To evaluate the benefits of this architecture, I will assess the data collected by the OPQ network at the University of Hawaii in order to determine if the claimed benefits of the architecture are observed in practice. The description of the evaluation processes for the benefits described in Section 1.3 are summarized below.

#### 1. Bandwidth usage is minimized: Bandwidth consumption of the OPQ system will be care-

fully monitored, recorded and compared to the bandwidth required to transmit the equivalent amount of raw data. A more detailed description of this evaluation can be found in Section 4.2.1.

- 2. Effects of latency are minimized: Latency limits of the triggering system will be examined. Since these are heavily dependent on the raw data storage ability of the OPQ Box, latency effects will be tested under various amount of memory allocated for this task. A more detailed description of this evaluation can be found in Section 4.2.3.
- 3. **Temporal locality:** Data acquired from the UH building level meters will be compared with the data acquired via the Napali triggering framework. This will allow me to establish the rate of false positives and negatives and evaluate the temporal locality triggering algorithm. Furthermore, the single location events, which are normally rejected by Napali will also be examined and compared to the data provided by the building power meters. A more detailed description of this evaluation can be found in Section 4.2.4.
- 4. Sink processing requirements are minimized: Synthetic benchmarks will be carried out on the sink node to determine the scalability of the triggering system. These scalability metrics will be compared with a synthetic benchmark of running multiple copies of the OPQ Box analysis software on the same node. This will allow me to compare the scalability of the sink node in the case of sending the entirety of raw data stream versus the Napali frameworks approach of only sending extracted metrics. A more detailed description of this evaluation can be found in Section 4.2.2.
- 5. Sub-threshold data acquisition: Evaluation of the sub-threshold data acquisition will performed in two ways. First, the triggering stream from the OPQ Boxes will be stored along with the acquired raw data. The triggering stream can be used to compute which fraction of devices would have self triggered if operating autonomously. Next the building level meters self triggering capabilities, will be compared to the Napali, sub-threshold triggering in order to compare Napali to a commercially deployed system. A more detailed description of this evaluation can be found in Section 4.2.5.

Although power failure resiliency and flexible privacy are claimed benefits of the Napali architecture, they will not be evaluated as part of this thesis research. Flexible privacy would require a much larger deployment, and a user study, which is beyond the scope of this project. Furthermore the power failure resilience of the Napali framework would require a significant development effort in development of the battery management system. Since complete power failures are quite rare, there is no guarantee that a single power outage would occur at the UH campus during the deployment.

The last claim of this dissertation is that the Napali architecture can be applied to other domains. Once the Napali framework is fully characterized, and its strengths and weaknesses are well understood, I intend to perform an in-depth literature review of other domains which could benefit from Napali-like approach to event detection. I will also characterize the kinds of design changes to existing sensors that the Napali Framework will require in order to apply it to these domains.

## CHAPTER 2 RELATED WORK

## 2.1 Edge computing

Projections performed by Forbes suggest that by 2025, more then 75 billion IOT devices will be connected to the Internet.[1] As the devices at the edge become more computationally capable and more numerous, it becomes imperative to share the computational load not only across the cloud services, but across the devices themselves. Furthermore, a large portion of these devices, such as home automation, do not require a connection to the cloud in the first place. Instead they require a connection to the edge IOT hub, or need the cloud service only to establish or broker communication with another IOT device. Edge computing is a subset of IOT research, which concerns itself with distributing the computational load across the devices at the edge of the of the network. [21]



Figure 2.1: Projected number of IOT devices worldwide.[1]

This change in computational strategy may seem inconsequential at first. However, upon deeper reflection it becomes clear that this is a major paradigm shift which brings IOT closer to the sensor network world it is often compared with. While pioneering work in the IOT always assumed a one or two way communication between the IOT device and the cloud service, utilizing TCP/IP as an end-to-end protocol, it is becoming clear that this approach is unsustainable, and is often undesirable. This communication model has clear disadvantages in wasted communication, computation, and privacy. Furthermore, the rigid computation communication model is not flexible enough to support devices which are beyond the edge of the TCP/IP network without an AD-HOC routing. [6]

The first attempt to address the bandwidth and latency issues arising with the widespread IOT adoption came in form of content delivery networks (CDNs).[6] CDNs circumvent the generic cloud information delivery problems by placing transparent caches geographically spread across the application domain as shown in figure 2.2. When a user or an IOT device makes a request for on object, this request is forwarded to the nearest CDN node for processing. If the node contains the object in its cache it is immediately forwarded to the requestee. Otherwise a request for the object is forwarded to the centralized cloud data store, and returned to the requesting device, as well as placed in the local cache. This approach has the advantage of moving the data closer to the end user, thus reducing latency, and taking advantage of the geographical locality. Another advantage of this method is the added resiliency of the CDN architecture to a single point failure. If a local cache node fails, its userbase, can be forwarded to another node, although incurring additional latency. Additionally, if a centralized data store becomes unreachable, the local cache nodes can to some extent mask it's outage by forwarding the data available locally. This approach does have some drawbacks. While it makes it easier to enable faster transactions regarding data, it is not trivial to move application logic to the local cache nodes. Furthermore, CDN methodology, still relies on a central mediator for device communication, even if the devices are located in the same room.



Figure 2.2: Content delivery network architecture. As described in the Google pattent.[6]

In an attempt to support a more diverse IOT ecosystem, current research is focused on moving the cloud service ever closer to the edge of the network. Since the majority of the IOT devices are located withing one hop of the Internet, the next logical place to locate a content provider is at the wireless basestation.[21] These servers, commonly referred to as cloudlets or fog servers, are collocated with various wireless basestations, which allows them to provide a location specific service to the user without using the Internet. This approach also provides uniform access and simplifies intercommunication between a variety of devices, including those that don't use TCP/IP. Unlike the localized cloud cache approach relied on by the CDNs, fog servers are built with the notion of moving not only data but also the application logic to the edge of the network. To facilitate inter-device communication between the devices using differing wireless protocols, fog servers can no longer rely on TCP/IP routing. Instead, TCP/IP becomes yet another transfer protocol along with Bluetooth, Zigbee, 3g etc, with routing between the devices implemented as a software service.[25] A few use-cases of such technology are already found in industry. Examples of these include airline/bus in-flight entertainment, and shopping mall directory apps. A block diagram of this infrastructure is shown in figure 2.3. In the future, emerging technologies which are sensitive to latency, such as virtual and augmented reality will benefit from fog computing, since it's inherently lower latency then the cloud counterparts.



Figure 2.3: Fog computing use in transportation. The bus cloudlet provides a cache for common data such as commuter schedules and traffic information, while routing other queries to the Internet.[25]

With the architecture for low latency communication between the edge devices and the fog provider established, the state of the art in edge computing research is focused on intelligent sharing of the computational resources in the fog system. Edge servers generally have a few orders more of computational capacity then the edge device, however, they service many such devices. Additionally, due to the fickle nature of radio links connecting the edge device to the fog service, the work sharing protocol must be able able to cope with link and packet loss. Finally, in the case of battery powered devices, the energy cost of transmitting the computational job, and receiving its result may exceed the cost of performing the computation locally. Finally, with mobile edge nodes, such as smart phones, and smart cars, computational offloading algorithms must be able to handle constant network reconfigurations as the edge nodes enter and leave the fog server geographical area. A number of algorithms have been proposed for efficient and robust computational sharing in fog environments. [17] [28] [29]

Napali fits in between the CDN and Fog server architectures. The power quality disturbances are generally localized to a specific area, so a sink placement which covers a small geographical area are preferred in order to reduce latency and reduce unnecessary communication with the centralized cloud location. Furthermore, sink driven measurement rate allows the OPQ Box to dynamically scale the computational and communication overhead. Finally the event, classification and analysis are similar to the computational offloading strategies currently under development in the edge computing field.

## 2.2 Distributed Power Quality Monitoring

Power quality monitoring is a long established field in the smart grid domain. However, the vast majority of research so far has focused on single point power quality monitoring.[23] Such research has extremely useful applications in industry, since it allows one to ascertain the absolute quality of the delivered power at a given location. However, since the power quality disturbances originate both from local sources and gridwide disturbances, single point monitoring is not particularly useful for smart grid research. Several projects have developed a distributed approach to power quality monitoring, the most prominent being the FNet project and the Power Standard Lab PQube deployment.

The FNet project designed, manufactured and deployed a Phase Measuremnt Unit (PMU), across over 300 locations across the united states.[32] PMU devices plug into an outlet, and sample the power line voltage at the rate of about 1.5 kS/s. The sampling is disciplined by GPS, and as such FNet devices are extremely sensitive to voltage frequency and phase angle. The precision of the FNet devices is 0.5mHz for frequency and  $0.02^{\circ}$  for phase angle. Collected data is sent to the collection service at 100ms intervals via the Ethernet connection. Using these devices FNet was able to observe several large power disturbances in the US power grid. The robustness and sensitivity afforded by the GPS receiver makes this project an excellent source of frequency data across large geographical area, however, the sampling rate of 25 samples per grid cycle is far too low to properly sample fast transients and sags. Furthermore, FNet provides no methodology for acquiring raw data for event disturbances which it records.

Power Standards Lab (PSL) has been an industry leader in power quality monitoring, and has authored several standards on the topic. Furthermore, PSL has developed and deployed a large number of power quality monitors called PQube across the world. The exact number of deployed devices is uncertain since a lot of the devices are deployed industrially and are not available to the public. However PSL has several publicly available devices, as well as several PQ datasets accessible for smart grid researchers. PQube devices are an industry standard for power quality monitoring, sampling at 12.8kSps for both the voltage and current waveforms.[2] Each PQube device is supplied with a NIST certificate of compliance and complies with the IEC 61000-4-7:2002 standard for PQ measurements. Incidentally, this standard was authored in part by the PSL staff. Similarly to the FNet PMU, PQube devices are GPS disciplined, additionally the sampling is phaselocked to the voltage waveform allowing for an even more precise metric extraction. Finally each PQube device is configurable with custom thresholds which allow it to record raw PQ event data for the location it's monitoring. PQube offers a centralized data collection option with flexible communication schemes ranging from Ethernet to Cellular. Since PQube devices monitor current in conjunction to voltage, its installation requires it to be placed into the electrical box of the target, by a licensed electrician.[27] Furthermore, the GPS synchronization requires addition of extra conduit to the electrical box to allow for an antenna. Finally, since the PQube devices are designed for single location measurements, distributed event detection using the PQube network is particularly difficult, with a lot of low magnitude gridwide events being incomplete or missing.

Unlike the single point monitoring solutions, the Napali framework is incapable of operating as standalone PQ monitor without a cloud sink. Furthermore, even with the event detection sink, the goal of Napali is to reject local anomalies in order to reduce the communication and computational overhead. While not as sensitive as the PQube device, the deployment price per unit is two orders of magnitude lower, while providing better sensitivity then the FNet PMU alternative, when running with GPS. The ability of operating using NTP, with WIFI connectivity, means that the OPQ Box deployment is much simpler then the FNet and PSL offering, without requirements of a clear view of the sky or additional wiring for Ethernet and GPS antennas. Finally the Napali distributed event detection system allows for acquisition of the entirety of the PQ disturbances including in locations where the disturbance has been greatly attenuated by the electrical distance. Thus, Napali is able to provide a more complete picture of the disturbance propagation throughout the smart grid.

## 2.3 Anomaly detection in Power Quality Monitoring Networks.

Anomaly detection in PQ domain remains an active topic of research. The goal of PQ event detection is to isolate the temporal regions where the voltage or current waveform deviates from the nominal by a given threshold. In some cases the aim is simply to notify a higher level control system in realtime manner that a disturbance is taking place. In other cases, the goal is to acquire the raw disturbance data for off-line analysis. Most of the detection methods rely on statistics and thresholding in order to detect PQ disturbances. Most of the literature concerns itself with single point detection, for purposes of protection of equipment downstream.[7][10] [22] Distributed power quality projects will generally utilize single single point detection across multiple devices in order to reconstruct gridwide propagation.[27]

With a wide deployment of smart meters, PQ researchers gain access to a networked platform which is perfectly positioned for PQ monitoring.[8] The major issue for smart meter real time monitoring is the bandwidth constraints. Smart meter deployment is envisioned to communicate via a mesh network with a stationary or mobile base station used for data aggregation. As such the bandwidth and connectivity is limited, thus requiring methodology which is capable of event detection in such environments. Generalized local likelihood ratio detection is one method for overcoming these limitations. This approach requires only a single bit to be forwarded from each smart meter indicating whether a disturbance is taking place or not. These bits are aggregated at the "master" meter, and if their sum exceeds a threshold a higher level control system is notified of the ongoing disturbance.[13] This approach is resilient to bandwidth limitations, and communication instability, however tunning thresholds for each individual meter requires a significant manual effort.

Systems designed for distributed PQ event detection using custom meters are prevalent in literature. A study at CERN utilized PQube devices with gapless recording which were later analyzed off-line, in order to ascertain the propagation mechanics of PQ events.[9] In a realtime domain Shang Li and Xiaodong Wang extended their work in [13] from smart meters to standalone devices, again advocating for single bit statistical based triggering generated by asynchronous meters.[12] Unfortunately their work has never been verified beyond simulation. The Transimeter project utilized an analog hardware event detector comprised of a high pass filter and a comparator for transient detection. These devices had two data paths for the voltage waveform, one to the National Instruments DAC board, one to the hardware trigger circuit. If a trigger circuit detected a transient, a flag was set on the NI DAC, which would in turn instruct the connected PC to send the data to the central server.[5] Unfortunately, the lack of cooperative detection and an inflexible trigger circuit makes this approach unappealing for modern power quality monitoring. Some of the more exciting work in PQ detection is modeling the most efficient placement of PQ meters in order to provide complete coverage for the power grid. [30] Another is using distributed detection for localization of the event source.[18] [20]

Napali differs from the smart-meter approach in the use of WiFi for communication, which greatly improves the communication constraints of the system. Since the OPQ Box is always connected to the power grid it is monitoring, power concerns are minimal. This allows Napali to implement more robust computational and communication strategies, not commonly possible with smart meter PQ monitoring. Since the triggering stream is generated in software, it is possible to switch the detection metrics without redesigning new hardware. Napali combines both cooperative PQ event detection and PQ event acquisition which makes it useful for future PQ event localization, and propagation research. [18] [20]

## CHAPTER 3 OPEN POWER QUALITY



Figure 3.1: Block diagram of the OPQ Power Quality monitoring system.

Open Power Quality (OPQ) power quality monitoring network utilizes residential power quality meters, called OPQ Boxes, in order to detect anomalies in the electricity distribution across the Oahu power grid. In addition to OPQ Boxes, the OPQ project utilizes cloud-based aggregation services for power quality event detection, classification and display. The block diagram of the OPQ network is shown in Figure 3.1.

The major components of OPQ are:

- OPQ Box: an in-house designed, open source power quality meter that conforms to Napali Framework requirements for the "source".
- Makai: data aggregation and event detection service that conforms to the Napali Framework requirements for the "sink"
- Mauka: event analysis and classification service.

The following sections describe the OPQ network components, services and protocols.

## 3.1 OPQ Box

OPQ Box is an in-house designed power quality meter which focuses on providing the means for cheap, extensible and accurate residential power quality measurements. The block diagram of the current revision of OPQ Box, OPQ Box2 is shown in the Figure 3.2a. A complete device in an acrylic enclosure is shown in Figure 3.2b.

The power system of the OPQ box2 electrically isolates most of the device from the AC mains power. An isolated AC-DC converter generates  $5V_{dc}$  from the mains  $120V_{ac}$ . 5V is used to power the Raspberry Pi, equipment connected to the expansion port, 3.3V regulators and voltage reference and an isolated DC/DC converter. 3.3V is used to power the isolated end of the isolation amplifier as well as the STM32F3 analog to digital converter/digital signal processor (ADC/DSP). The hot



(a) OPQ Box2 Block Diagram. The power path is in red, signal path is in green and the digital IO is in black.



(b) OPQ Box2 in an acrylic enclosure.

Figure 3.2: (a) OPQ Box2 block diagram and (b) production OPQ box ready for deployment

side of the isolation amplifier is powered from the isolated DC/DC converter. This allows OPQ box to function with the battery attached to the expansion port, so that it may retain data and continue to operate during a power outage.

The analog signal path of the OPQ Box2 is complicated by the fact that the STM32F3 ADC/DSP is electrically isolated from the mains power. A previous iteration of the OPQ Box, OPQ Box1, overcame this by employing small circuit board mount isolation transformer. Unfortunately it was found that the frequency response of these transformers varied wildly between individuals, thus incurring a lengthy calibration process for each device. Design on the OPQ Box2 solved this issue by using an AMC1100 isolation amplifier as the isolation component. Internally AMC1100 consists of a single die comprised of a  $\Sigma\Delta$  analog to digital and digital to analog converters. These converters are separated by a silicon glass region on the integrated circuit which acts as a coupling capacitor. Since the signal passes the isolation boundary as a  $\Sigma\Delta$  encoded digital signal, it incurs no distortion and no additional calibration is required. In order to match the dynamic range of the AMC1100 the  $120V_{ac}$  is passed through the resistor divider to attenuate it to  $120mV_{ac}$ . The input and output of the isolation amplifier is filtered with a passive first order anti-aliasing filter. Isolated signal is then digitized via a 16bit ADC of the STM32F3 DSP at 12KSps, which gives 200 data samples per grid cycle. Internally digitization process runs asynchronously with the respect to the the DSP CPU, in order to minimize timing jitter. It was verified that the sampling jitter of the ADC is less then 1us, however due to limited precision of equipment an exact figure was not established. Data stream in its digital form is transferred to the Raspberry Pi single board computer (SBC) for analysis.

Raspberry Pi SBC is responsible for signal analysis and anomaly detection. The Raspberry Pi model used in OPQ Box is the Pi Zero W equipped with 256MB of main memory and a single core 1GHz ARM11 CPU. Furthermore, Pi Zero W is equipped with an on-board 802.11n



Figure 3.3: Block diagram of the OPQ Box 2 software stack.

WIFI transceiver, which removes the need for an external WIFI dongle used in previous OPQ Box devices. The software stack of the Raspberry Pi aims to deliver a full featured power quality analysis framework despite its rather limited hardware capabilities. A block diagram of the software stack is shown in Figure 3.3. Digital data is transferred from the DSP to the Raspberry Pi via Serial Peripheral Interface, with the Pi acting as the master and the DSP as a slave device. A hardware interrupt line is used to inform Pi software that the DSP is ready for the data transfer. During the initial design of the OPQ box 2 software, SPI data transfer was attempted in userland. However due to the lack of support for DMA in the SPI kernel-to-userland bridge, a large portion of the CPU time was spent facilitating data transfer, resulting in degraded analysis performance as well as missed data samples. Current revision of the OPQ Box 2 software stack utilizes a kernel driver for management of SPI bus. Internally OPQ driver maintains a ring buffer of 16 windows each 200 data samples in size. Upon the receiving the interrupt for the DSP, the CPU sets up the DMA transfer and the DMA engine transfers a 200 sample window into the kernel memory without CPU interaction. This scheme requires the CPU to only service 60 interrupts a second, with each interrupt requiring on the order of 100 instructions, yielding the CPU utilization of less then 1% in normal operation. Userland applications communicate with the kernel driver using a file descriptor, where every *write* system call yields 200 samples of raw waveform. Thus the smallest window that a userland application will process is a single AC cycle of the grid mains.

Userland component of the OPQ Box 2 software is a multi-threaded extensible analysis framework called Triggering. The reader thread is responsible for transferring and accumulating data from the kernel driver. The smallest data buffer that the Triggering application processes at any given time is 10 grid cycles or 2k samples. Once the cycles are transferred to the userland and timestamped, they are passed to the analysis thread for feature extraction, as well as to the Raw Data Ring Buffer (RDRB). Since internally all data is addressed using shared pointers, during data duplication no copying is required. RDRS is capable of buffering up to an hour of historic data before it's overwritten resulting in the RDBS maximum size of 100MB.

Analysis thread of the Triggering application performs feature extraction of the raw data windows of 2000 samples. At the moment only two metrics are calculated:

• Fundamental frequency.

• RMS Voltage.

Fundamental frequency is calculated by computing the zero crossings of the AC waveform. Since a sinusoid will have two zero crossings for a full cycle the frequency can be calculated as:

f



$$=\frac{1}{\frac{2}{n}\sum\limits_{k=0}^{k=n}\Delta t_{k}}$$
(3.1)

Figure 3.4: Filters used for mains frequency calculation. (a) Downsampling filter gain. (b) Downsampling filter impulse response. (c) Lowpass filter gain. (d) Lowpass filter impulse response.

Where the  $\Delta t_k$  is the k'th time between two adjacent zero crossings. In order to improve the accuracy of the frequency calculation one must first filter out as much of out of phase noise as possible. Since our sampling rate is quite high (12kSps) and the fundamental frequency is quite low (60Hz) it is very computationally expensive to perform this filtering in a single step. Instead filtering is accomplished via a set of two finite impulse response (FIR) filters shown in Figure 3.4b and 3.4d. First the Down sampling filter is applied to the raw waveform, which results in the frequency response shown in Figure 3.4a. As is evident by the plot the frequency content of

the result is 0-600Hz, Thus it can be downsampled to the  $\frac{N}{10}$ , or 200 samples without aliasing. Next the low pass filter is applied to the downsampled waveform with the frequency response shown in Figure 3.4c.This resulting frequency content is 0-100Hz, thus all of the higher frequency harmonics and noise are removed. Finally the twice filtered downsampled waveform is used to estimate the fundamental frequency according to the Equation 3.1. The zero crossings themselves were calculated by using linear interpolation between two points which bracket the time axis.

In order to verify the error in our frequency measurement, a function generator (SIGLENT SDG1025) was used to inject a 60Hz  $120mV_{pp}$  superimposed with 1% white noise into the input of the AMC1100 anti-aliasing filter, bypassing the voltage divider. The resulting frequencies were calculated and recorded by the OPQ Box 2 and histogramed as shown in Figure 3.5a. It was found that the OPQ Box 2 overestimated the frequency by  $300\mu Hz$  with  $\sigma = 230\mu Hz$ . All electrical generation systems connected to the grid run synchronously with each other, meaning that while small variations in voltage are common across locations, the fundamental frequency and phase must remain strictly in sync. This effect is demonstrated in Figure 3.5b, which is a frequency fluctuation event recorded on November 8 2017. While the two devices were separated by ten miles, their frequency measurements track closely together.



Figure 3.5: Filters used for mains frequency calculation. (a) Downsampling filter gain. (b) Downsampling filter impulse response. (c) Lowpass filter gain. (d) Lowpass filter impulse response.

Root mean square voltage  $(V_{rms})$  in electrical power is the equivalent value of DC voltage which would dissipate the same power in the resistive load.  $V_{rms}$  is a convenient measure for detecting voltage sags and swells, since they result in nominally higher and lower computed value. For the sinusoidal signal  $V_{rms}$  can be calculated from the peak to peak value  $(V_{pp})$  as:

$$V_{rms} = \frac{V_{pp}}{2\sqrt{2}} \tag{3.2}$$

It is common for multimeter to employ the equation above for computing  $V_{rms}$ . However this equation will only work for a perfect sinusoid, and thus does not result in a suitable metric for identifying power quality disturbances. Instead OPQ Box 2 computes  $V_{rms}$  as follows:

$$V_{rms} = \sqrt{\frac{1}{n} \sum_{k=0}^{k=n} V_k^2}$$
(3.3)

Similarly to the frequency calculation OPQ Box 2 will use a 10 cycle window for a single  $V_{rms}$  calculation, however unlike the frequency calculation the input is not filtered a priori. An example of a power quality disturbance which exhibits low  $V_{rms}$  is shown in Figure 3.6a and 3.6b. These disturbances are the result of a lighting strike recorded by two OPQ Box 2 devices on Nov 1, 2017.



Figure 3.6: A lightning strike recorded by two OPQ Box 2 devices separated by 10 miles. (a) A lightning strike manifested as  $V_{rms}$  dip which lated 11 cycles. (b) As a consequence of using NTP these devices have  $\frac{1}{2}$  cycle mismatch in reported timestamps.

Computed fundamental frequency and  $V_{rms}$  are transmitted to the Makai service for aggregation. Data transmission is handled using 0MQ software stack with Curve25519 elliptic curve encryption. Each device holds a unique of private and public keys, as well as the servers public key, allowing both the Makai service and the OPQ Box 2 to verify it's peer. Internally metrics transmission uses 0MQ's PUB/SUB protocol. This protocol is a publish subscribe one to many topology, with each message containing the topic, and a payload. Additionally 0MQ pub-sub topology allows for multiple sub peers with subscriptions forwarder to the publisher automatically via a side channel. This allows for the aggregation service to be spread across multiple nodes, with minimal network overhead.

If the aggregation service determines that an anomaly has occurred, it is able to request raw waveform from the OPQ Box 2 RDRB via a separate 0MQ pub sub channel. If the RDRB buffer contains data for the requested temporal range, OPQ Box 2 will send the available data to the aggregation service via a push pull 0MQ channel. Protobuf message serialization is used to encode messages across the OPQ ecosystem.

In order to make a distributed measurement, all of the OPQ Boxes on the OPQ network need to maintain an accurate time reference. Time synchronization across multiple OPQ Boxes is accomplished using the Network Time Protocol. The expansion port of the OPQ Box 2 supports a GPS receiver, however using it is detrimental to the OPQ Box 2 utility. GPS receivers require line of sight to the sky, and since the with out on-board real-time clock, every power interruption requires a GPS cold start. NTP performance has been verified against GPS resulting in time error of  $8ms \pm 5ms$  which is typical for NTP running over the Internet with a close by NTP server. This error is visualized in a Figure 3.6b. With a large coincidental  $V_{pp}$  drop across two devices, a 7ms phase error is clearly visible.

During the evaluation phase of the OPQ Box a new metric based on wavelet decomposition will be developed for detection of transients. This will allow the OPQ Box to detect slow voltage fluctuations using  $V_{rms}$ , and fast changes using the new wavelet metric.

## 3.2 OPQ Makai

OPQ Makai implements the Napali Framework requirements for a "sink". It is a distributed extensible microservice framework responsible for receiving the triggering stream from the OPQ Boxes, locating anomalous temporal regions and requesting raw waveform for the anomalous time ranges. As evident from the block diagram shown in Figure 3.7, Makai consists of three major components: Acquisition Broker, Triggering Broker, and the Acquisition Service.

#### 3.2.1 Triggering Broker

The triggering Broker is perhaps the simplest component of the OPQ Makai system. The triggering stream generated by the OPQ Boxes is encrypted to preserve users privacy. In order to minimize the CPU time spent decrypting the data across multiple OPQ services, the Triggering Broker is used to decrypt the data and send clear text measurements across the rest of the OPQ ecosystem. Triggering Broker uses the 0mq sub socket to receive data form OPQ Boxes, and send, it via a pup socket to any connected client. Another potential role of the Triggering Broker is load balancing. Since the 0MQ pub/sub sockets are subscription based, services connecting to the triggering broker can selectively receive measurements only from OPQ Boxes they are interested in.



Figure 3.7: Block diagram of the OPQ Makai.

## 3.2.2 Acquisition Broker

The Acquisiton Broker manages the two way communication between the OPQ Boxes and the rest of the cloud infrastructure. Unlike the triggering stream which originates from the OPQ Box, two way communication is always initiated by the cloud services. Two way communication is realized via a command response interface, where the OPQ service initiates the communication by sending in clear text command to the Acquisition Broker, which then forwards it in the encrypted form to the appropriate OPQ Boxes. There are three command types which can be handled by the Acquisition Broker:

- (PING) Ping: The ping command is broadcast periodically across all of the OPQ Boxes, in order to monitor the health of the OPQ network. The OPQ Box response to this command contains diagnostic information, such as the timestamp of the last event requested, ip address and the name and strength of the WIFI network the OPQ Box is connected to.
- (CMW) Change measurement window: This command allows to vary how often a triggering stream message is generated and delivered to the triggering broker. This is accomplished by varying the length of the temporal window used to derive the triggering metrics. This allows the OPQ system to analyze a finer grained measurements if a potential anomaly is taking place.
- (RD) Send raw data: This command instructs the OPQ Box to send data from the its raw data buffer to the cloud.

The last command in particular is unique, because its response is trapped by the Acquisition Broker, as opposed to being passed to the agent which initiated communication. In this case the raw data is serialized and stored in the Mongo database. Furthermore, the Acquisition Broker notifies all connected agents that a new anomaly has been recorded via a 0mq pub message.

## 3.2.3 Acquisition Service

The Acquisition Service middleware resides in between the Triggering and Acquisition Brokers. It processes the triggering stream and manipulates it using the **CMW** command and if an anomaly is suspected requests raw data from boxes via the **RD** command. The Acquisition Service app itself does not perform any analysis of the triggering stream. Instead it provides a loadable plugin interface which allows for a runtime hotplugable analysis. Since the OPQ Makai core is written in the Rust programing language, the loadable object interface supports supports any programing language with LLVM bindings. As such OPQ Makai plugins can be developed in C/C++, Rust, or even Lisp. Currently three plugins are implemented:

- Print Plugin: Prints the triggering stream to the stdout for debugging purposes.
- **Ketos:** A Lisp interpreter built on top of the LLVM for manipulating/debugging the triggering stream.
- **Threshold:** This plugin monitors the triggering stream, while looking for metrics which exceed the user defined thresholds. If the threshold is crossed, data from all of the OPQ Boxes is requested for the offending time interval via the RD command.

Future revisions of Acquisition Service will likely include a Python interpreter plugin, however due to the persistent global state in the Python interpreter only a single Python instance can run in the given address space.

## 3.3 OPQ Mauka

OPQ Mauka service is responsible for higher level classification and filtering of the anomalies generated by the OPQ Makai. Since anomalies generation only relies on the triggering stream features and not raw data, OPQ Makai is not able to ascertain if the anomaly is an actual power quality event, event type, or its severity. OPQ Mauka on the other hand operates on the raw data, thus it is able to perform high level analysis to meet industry standards for vent classification. The block diagram of the OPQ Mauka is shown in 3.8.

Currently OPQ Mauka supports the following classification strategies:

- **ITIC** Power acceptability curve used to classify short term voltage sags.
- IEEE 1159 Voltage Voltage classification based on the IEEE 1159 power quality standard.
- Brownout Detection Classification of medium to long term voltage sags.
- Total Harmonic distortion Classification of events via harmonic analysis.



Figure 3.8: Block diagram of the OPQ Mauka.

Once the anomaly is classified by OPQ Mauka, and the power quality characteristics are confirmed, it may be aggregated with other anomalies to form a disturbance. Disturbances are composed of raw box data, analysis results as well as expert annotations and other metadata.

## 3.4 OPQ View

OPQ View is the primary user interface to the OPQ ecosystem. View is written in JavaScript using the Meteor framework, and provides a robust and easy to use interface to the OPQ Box triggering stream, Makai triggering anomalies, and to the Mauka PQ disturbances. Furthermore, View provides an administration interface for initial setup and maintenance of the OPQ devices, and services. Finally OPQ View monitors the health of the OPQ components, keeping track of the individual box uptimes, and component failures. A screenshot of the recent OPQ View build is shown in Figure 3.9



Figure 3.9: Screenshot of a recent OPQ View build.

## CHAPTER 4 EVALUATION

Validation of Napali hinges on the validation of performance of the OPQ Boxes. Some of the detection capability of the OPQ Boxes has already been characterized as shown in the previous section, however transient response and the  $V_{rms}$  response are yet to be confirmed. Once the OPQ Box sensitivity has been characterized, the detection capability of Makai system must be crosschecked. First synthetic triggering streams will be injected into Makai. Next the full system will be deployed at the University of Hawaii Manoa for in-situ validation.

## 4.1 OPQ Box Validation.

So far only the frequency sensitivity of the OPQ Box has been validated empirically. The transient and as well as  $V_{rms}$  response characterization will first be performed using synthetic data. Robust methods for generating power quality events are present in literature, and thus no new research for single device validation is required.[11][26] Next once the DSP software is characterized, synthetic PQ data will be loaded into a SDG1025 signal generator, and fed into the OPQ Box hardware. Any discrepancy between the DSP, and hardware-in-the-loop characterization will be noted and analyzed. The main characteristics to be validated are as follows:

- DC response: Accuracy of the OPQ Box in measurement of a DC voltage.
- $V_{rms}$  response: Accuracy of the OPQ Box in measurement of a small changes in the amplitude of AC waveform.
- THD response: Accuracy of the OPQ Box in measurements of small harmonics mixed with a large fundamental AC waveform.
- Transient response: Validating the response of the OPQ Box metrics to various transients.

These tests will provide a baseline for the detection capabilities of the Napali system, and result in a publication regarding the OPQ Box detection capabilities.

## 4.2 Napali Validation.

In order to validate the OPQ system as a whole, it will be deployed across the University of Hawaii Manoa campus (UH). This location is ideal because it is a relatively isolated microgrid connected to the Oahu powergrid only via a single 46kV feeder as shown in Figure 4.1. Another advantage of the UH campus is the high number of smart meters deployed across various levels of the power delivery infrastructure. While these meters are mainly geared towards monitoring power consumption they do have some power quality monitoring capabilities. Data provided from these meters can be used in two distinct applications. First of all, this data can be used to pinpoint the sections of the University of Hawaii power grid which experience a higher likelihood of power quality disturbances. These portions of the grid will have a higher spacial density of OPQ Boxes. Secondly, data from the campus deployed meters can be used as ground truth for comparison against the measurements, and analysis performed by the OPQ project. The location of smart meters in the grid topology is shown in figure 4.1 as the M nodes. As evident by the meter location none of them are monitoring the consumer level power and mainly focus on the higher voltage power delivery. This placement is evidenced from the smart meter role as a consumption monitor, and thus the deployment of the OPQ Boxes at the residential level will compliment the current power quality monitoring capabilities without introducing redundancies. Finally, the University of Hawaii power grid is supplying a highly diverse infrastructure. Beyond the traditional residential equipment such as computers and consumer grade electronics, the UH power grid powers scientific and laboratory equipment, machine shops, and server farms. All of these elements have varying requirements/tolerances for power quality anomalies as well as different levels of power quality "pollution". Furthermore, some of the electricity consumers in the UH campus are entirely unique. For example, the free electron laser located in the Watanabe Hall is one of the only free electron lasers in the world, and the impact/sensitivity of power quality on the instrument are completely unstudied.



Figure 4.1: University of Hawaii at Manoa power delivery infrastructure.

There are 74 smart meters deployed across the UH campus. These meters measure the fundamental frequency  $V_{rms}$ , power consumption, reactive power, and power factor. Data from these meters will be cross-referenced with the Napali detection system in order to ascertain it's benefits. Validations of the benefits of the Napali framework will follow the framework described in Section 1.6. Here we examine the analysis required to complete each claim in detail.

#### 4.2.1 Napali Bandwidth usage

In order to analytically compute the bandwidth savings of the Napali compared to a system which sends all the data to the sink, I will keep track of the amount of bandwidth the Napali monitoring system consumed during its deployment on UH campus. This will in turn be compared to the bandwidth required by the OPQ Boxes if they were sending the entirety of the data to the sink, and establish the bandwidth efficiency of the Napali framework. In order to monitor the bandwidth of the system in-situ, an iptables bandwidth accounting will be enabled on the sink node. With the accounting enabled, I will be able to extract the exact amount of bytes transfered by the Napali framework during the deployment period. Since the sampling rate of the OPQ Box is well characterized, and the number of OPQ Boxes is fixed, it is trivial to calculate the amount of raw data generated by the OPQ network during any time period. In order to make this comparison fair, the raw data bandwidth will be scaled by the compression ratio of the state of the art compression algorithm specifically designed for power quality measurements.[31]

Operating at 12kSps, OPQ Box produces raw data at 24KB/s. With state of the art compression operating at 90% compression ratio and a 5% overhead of TCP/IP and meta-data, one can expect a 3KB/s stream of raw data for each OPQ box if it were to send the entirety of it to the sink. In synthetic experiments, under steady state conditions, OPQ Box consumes 0.5KB/s of bandwidth sending metrics to the Makai. However, this figure does not take into account sending the raw data samples in case an event is detected. The result of bandwidth consumption evaluation will provide a comparison of the bandwidth consumption by Napali methodology of sending metrics and temporal raw data windows of interest compared to sending the entirety of raw data to the sink.

#### 4.2.2 Sink processing requirement under the Napali Framework

In order to evaluate the scalability of the sink node under the Napali framework, synthetic data containing a small amount of distributed events will be injected into the Makai triggering system. While the synthetic data for a single point PQ disturbance is easily generated, distributed PQ event generation is not well understood. However there is some literature concerning power fault propagation and localization. [18] [20] The main takeaway from these authors is the energy and amplitude of the event diminishes with electrical distance from the source. As such, by generating a single point PQ event as described in [11][26], and linearly scaling it based on the simulated electrical distance from the source, a distributed PQ event ensemble can be generated. These events will be ran through the simulated OPQ Box DSP stack and the extracted metrics will be propagated to the Makai aggregation sink. The number of simulated OPQ Box devices that can be supported on a single node will be recorded, and will provide the baseline for the Makai sink node capabilities. Next, the same dataset will be processed on the same node using simulated OPQ Box software stacks. The amount of concurrent processes which are able to keep up with the targeted

sampling rate of 12kS/s will be determined and recorded, and compared to the amount of devices which can be supported by Makai.

The result of this evaluation will provide a scalability metric for the Napai framework for a single sink node. There are independent metrics which will be presented. First, the number of OPQ Boxes which can reliably send metrics to a single Makai instance for processing will be established. This number is expected to be quite high since the bandwidth and computational requirements for processing the triggering stream are quite low. Second, the number of OPQ Boxes, that can experience and reliably send an event to Makai at the same time will be established. The bandwidth requirement for raw data acquisition is significantly higher then the triggering stream, however there is no computation involved in storing raw data frames. Both of these metrics will be compared to the scalability of performing the entirety of processing on the sink node.

#### 4.2.3 Effects of latency in the Napali framework

The latency of Napali triggering system has a significant impact on its ability to read out complete raw data events. Using generated distributed events as a baseline, I will be able to tune the threshold and temporal requirements for Makai detection algorithms. Furthermore, temporal, spacial, and amplitude noise will be injected into the generated datasets, to simulate various uncertainties with regards to data collection, such as local noise, and NTP offset errors. Taking into the account the detection latency of Makai, if some of the requested data is no longer available on the OPQ Box, only a partial time window will be returned. These events will be marked as incomplete and their fraction as compared to the total number of events recorded will be used to establish the latency tolerance of the Napali framework. Synthetic benchmarks will be carried out to establish the latency that the Napali system can incur without losing a portion of the event. Since this is highly dependent on the amount of storage allocated for the RDRB, these experiments will be carried out with various RDRB settings.

Because OPQ Boxes operate using public University of Hawaii WIFI, the latency figures for data transmission are expected to be very dynamic. In situations with large network contention, greater then 100ms one way packet latency can be expected. This latency is exacerbated since at least three separate communication steps are required before raw data is received by Makai. First the metric has to be sent to Makai, next if Makai detects an event, a data request needs to be sent to the affected boxes. Finally boxes will forward raw data to the Makai sink. I expect that with RDRB capacity of storing 5 seconds of raw data, no raw data events of 1s or shorter will be missed. This figure allows for 1 second of transfer latency, 3 seconds of Makai analysis latency, and leaves 1 second of data in RDRB for readout. However, these figures can only be validated in a real world situations.

#### 4.2.4 Temporal locality triggering of the Napali framework

Once the OPQ Box is fully validated and the Makai detection thresholds are tuned using synthetic datasets, the Napali system will be fully deployed at the University of Hawaii at Manoa. Every time the Napali detects an event, both OPQ Boxes and building meters will be queried for data. While it may be unfeasible to query raw data from the UH metering infrastructure, metrics are readily available. This data will be used to ascertain the proportion of false positive events detected by Napali. Additionally, the internal single point fault detection mechanism of the UH power meters will be used in conjunction with the events detected by Napali to measure the rate of false negative events. Both the false negative and false positive measurements will be used to ascertain the detection efficient of the Napali framework. This analysis will also include an evaluation of Napali's ability to reject single point anomalies. For a portion of OPQ deployment, every event triggered by a single device will be captured. These events will be analyzed in order to determine if a gridwide anomaly was incorrectly classified as a single point disturbance.

The goal of Napali is not to to provide a zero false positive rate. Once raw data is stored, higher level processing can further filter and classify it using more computationally expensive techniques. As long as the bandwidth consumption of Napali compares favorably to sending the entirety of raw data to the sink, any rate of false positives can be tolerated. False negatives on the other hand are the primary metric subject to optimization. Ideally a zero rate of false positives would be expected, however as with any real-world system I do not expect that to be the case. This evaluation will determine the triggering efficiency of the Napali framework when compared to the detection ability of a commercially available system.

#### 4.2.5 Sub-threshold Data Acquisition

The Napali methodology will be compared with the single point anomaly detection approach. In order to do that I will compare the extent to which sub-threshold events are missed by the UH metering infrastructure. In a large distributed event, if a portion of events are not detected by the UH meter's single point detection, but picked up by the Napali framework, these events will be flagged and analyzed for their merit. This will in turn provide a metric of distributed detection ability of the Napali framework compared to commercial system. Furthermore, for a portion of the deployment the triggering stream from the OPQ Boxes will be stored along with the acquired raw data. The triggering stream can be used to compute which fraction of devices would have self triggered if operating autonomously. This will provide the baseline for sub-threshold triggering efficiency of the Napali system, with respect to the single point detection ability of the OPQ Box.

This evaluation will compare Napali performance to the single point detection mechanisms currently in deployment. I expect that Napali will outperform these strategies, and provide a more complete picture of gridwide anomalies as they propagate through the UH power grid. It is possible that no sub-threshold events will be recorded during the UH deployment. UH campus is quite small, perhaps too small for an anomaly in one building not to impact the rest of campus. In this case the sub-threshold data acquisition will remain an open topic for future work and a larger geographical deployment. Regardless, as long as I am be able to validate the single point anomaly rejection ability of Napali as described in Section 4.2.4, I will be able to conclude that Napali has a distinct advantage over single point detection methods. Single point anomalies are not important in smart grid monitoring, since they originate from the consumers side of the meter, and should be ignored. In fact, recording these events may be detrimental to the privacy of the end-user, since it may give clues on their activities as shown in Figure 1.2 a and b. However since privacy implication of power quality monitoring are outside the scope of this project, this study will remain as a point of future work.

## 4.3 Napali Framework in Other Domains

The work outlined above will allow me to characterize the benefits and drawbacks of the Napali architecture. Armed with this knowledge I intend to examine the examine the possibility of using a similar architecture in other domains. The Napali approach could prove useful in general sensor network design, as well as IOT and other domains which rely on anomaly detection. This contribution will involve an additional literature review of anomaly detection systems, and proposals to integrate Napali-like methodology to improve their efficiency and reduce their bandwidth consumptions. It will also include a discussion of hardware design changes required to existing sensors in these other domains in order to support the Napali framework.

## 4.4 Schedule

т	<b>`</b>	1 .	•	1		•	1 • 1	•	. 1	•	1		1 1.				1 C
- H	<u> </u>	OTT 1	ic t	ho	timo	Ino 1	docorit	nn m	tho	malor	rocoarch	o of ivition	londing	110 f	· O Y	$m \tau \tau$	dotongo
	101		טכו	IIC -	UTTE-	ппс (	перени	лпе	UHC.	шаю	research	activities	leaung	11171	JU I	IIV	uerense.
								. 0		- J -				· · · · ·			

Activity	Start Date	End Date
OPQ Box 2 Validation	October 1 <sup>st</sup>	November $31^{st}$
Makai Validation	October 15 $1^{st}$	November $30^{\text{th}}$
Data Collection	December $1^{st}$	TBD
OPQ Box 2 Instrumentation paper	January $1^{st}$	February $1^{st}$
Data Analysis	February $1^{st}$	March $1^{st}$
Introduction Dissertation Chapter	March $1^{st}$	March $15^{\text{th}}$
Literature Review Dissertation Chapter	March $15^{\text{th}}$	April $1^{st}$
Experimental Design Dissertation Chapter	April $1^{st}$	April $15^{\text{th}}$
Analysis Dissertation Chapter	April $15^{\text{th}}$	May $15^{\text{th}}$
Conclusions Dissertation Chapter	May $15^{\text{th}}$	June $15^{\text{th}}$

## BIBLIOGRAPHY

- [1] Iot market predicted to double by 2021, reaching \$520b. Accessed: 2018-09-30.
- [2] Pqube specifications. Accessed: 2018-09-30.
- [3] Ieee recommended practice for monitoring electric power quality. IEEE Std 1159-2009 (Revision of IEEE Std 1159-1995), pages c1-81, June 2009.
- [4] Frede Blaabjerg, Remus Teodorescu, Marco Liserre, and Adrian V Timbus. Overview of control and grid synchronization for distributed power generation systems. *IEEE Transactions on industrial electronics*, 53(5):1398–1409, 2006.
- [5] P Daponte, M Di Penta, and G Mercurio. Transientmeter: A distributed measurement system for power quality monitoring. *IEEE Transactions on Power Delivery*, 19(2):456–463, 2004.
- [6] Joshua D Gagliardi and Timothy S Munger. Content delivery network, June 14 2011. US Patent 7,962,580.
- [7] Irene YH Gu, Nichlas Ernberg, Emmanouil Styvaktakis, and Math HJ Bollen. A statisticalbased sequential method for fast online detection of fault-induced voltage dips. *IEEE Trans*actions on Power Delivery, 19(2):497–504, 2004.
- [8] Joel Höglund, Dejan Ilic, Stamatis Karnouskos, Robert Sauter, and Per Goncalves Da Silva. Using a 6lowpan smart meter mesh network for event-driven monitoring of power quality. In Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on, pages 448–453. IEEE, 2012.
- [9] K Kahle. Power converters and power quality. arXiv preprint arXiv:1607.01556, 2016.
- [10] Masoud Karimi, Hossein Mokhtari, and M Reza Iravani. Wavelet based on-line disturbance detection for power quality applications. *IEEE Transactions on Power Delivery*, 15(4):1212– 1220, 2000.
- [11] V Kumar and M Kaur. Power quality event generation in matlab/simulink environment. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 4(7):6055–6059, 2015.
- [12] Shang Li and Xiaodong Wang. Monitoring disturbances in smart grids using distributed sequential change detection. In Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2013 IEEE 5th International Workshop on, pages 432–435. IEEE, 2013.

- [13] Shang Li and Xiaodong Wang. Cooperative change detection for voltage quality monitoring in smart grids. *IEEE Trans. Information Forensics and Security*, 11(1):86–99, 2016.
- [14] Yong Liu, Shutang You, Wenxuan Yao, Yi Cui, Ling Wu, Dao Zhou, Jiecheng Zhao, Hesen Liu, and Yilu Liu. A distribution level wide area monitoring system for the electric power grid-fnet/grideye. *IEEE Access*, 5:2329–2338, 2017.
- [15] Johan Morren, Sjoerd WH De Haan, Wil L Kling, and JA Ferreira. Wind turbines emulating inertia and supporting primary frequency control. *IEEE Transactions on power systems*, 21(1):433–434, 2006.
- [16] Mohd Izhwan Muhamad, Norman Mariun, and Mohd Amran Mohd Radzi. The effects of power quality to the industries. In *Research and Development*, 2007. SCOReD 2007. 5th Student Conference on, pages 1–4. IEEE, 2007.
- [17] Jessica Oueis, Emilio Calvanese Strinati, and Sergio Barbarossa. The fog balancing: Load distribution for small cell cloud computing. In Vehicular Technology Conference (VTC Spring), 2015 IEEE 81st, pages 1–6. IEEE, 2015.
- [18] Anthony C Parsons, W Mack Grady, Edward J Powers, and John C Soward. A direction finder for power quality disturbances based upon disturbance power and energy. In *Harmonics and Quality of Power Proceedings, 1998. Proceedings. 8th International Conference On*, volume 2, pages 693–699. IEEE, 1998.
- [19] S Peisert, R Gentz, J Boverhof, C McParland, S Engle, A Elbashandy, and D Gunter. Lbnl open power data. 2017.
- [20] Boštjan Polajžer, Gorazd Stumberger, and Sebastijan Seme. Evaluation of different methods for voltage sag source detection. *Energija*, 59(1-4):0–0, 2017.
- [21] Mahadev Satyanarayanan. The emergence of edge computing. Computer, 50(1):30–39, 2017.
- [22] Y-J Shin, Edward J Powers, Mack Grady, and Ari Arapostathis. Power quality indices for transient disturbances. *IEEE Transactions on Power Delivery*, 21(1):253–261, 2006.
- [23] Pedro LM Silva, José A Afonso, Vítor Duarte Fernandes Monteiro, JG Pinto, and João L Afonso. Development of a monitoring system for electrical energy consumption and power quality analysis. 2017.
- [24] S. Sivanagaraju. Electric Power Transmission and Distribution. Dorling Kindersley, 2008.
- [25] Xiang Sun and Nirwan Ansari. Edgeiot: Mobile edge computing for the internet of things. 54:22–29, 12 2016.

- [26] Rodney HG Tan and VK Ramachandaramurthy. Simulation of power quality events using simulink model. In *Power Engineering and Optimization Conference (PEOCO)*, 2013 IEEE 7th International, pages 277–281. IEEE, 2013.
- [27] Alexandra Von Meier, David Culler, Alex McEachern, and Reza Arghandeh. Microsynchrophasors for distribution systems. In *Innovative Smart Grid Technologies Conference* (*ISGT*), 2014 IEEE PES, pages 1–5. IEEE, 2014.
- [28] Kaiqiang Wang, Minwei Shen, Junguk Cho, Arijit Banerjee, Jacobus Van der Merwe, and Kirk Webb. Mobiscud: A fast moving personal cloud in the mobile network. In Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges, pages 19–24. ACM, 2015.
- [29] Shiqiang Wang, Guan-Hua Tu, Raghu Ganti, Ting He, Kin Leung, Howard Tripp, Katy Warr, and Murtaza Zafer. Mobile micro-cloud: Application classification, mapping, and deployment. In Proc. Annual Fall Meeting of ITA (AMITA), 2013.
- [30] Dong-Jun Won, Il-Yop Chung, Joong-Moon Kim, Seung-Il Moon, Jang-Cheol Seo, and Jong-Woong Choe. A new algorithm to locate power-quality event source with improved realization of distributed monitoring scheme. *IEEE Transactions on Power Delivery*, 21(3):1641–1647, 2006.
- [31] Dahai Zhang, Yanqiu Bi, and Jianguo Zhao. A new data compression algorithm for power quality online monitoring. In Sustainable Power Generation and Supply, 2009. SUPERGEN'09. International Conference on, pages 1–4. IEEE, 2009.
- [32] Yingchen Zhang, Penn Markham, Tao Xia, Lang Chen, Yanzhu Ye, Zhongyu Wu, Zhiyong Yuan, Lei Wang, Jason Bank, Jon Burgett, et al. Wide-area frequency monitoring network (fnet) architecture and applications. *IEEE Transactions on smart grid*, 1(2):159–167, 2010.