

# Localizing Loads in Microgrids Using High-Precision Voltage Phase

Maxim Buevich  
Carnegie Mellon University  
Pittsburgh, PA  
maximb@cmu.edu

Anthony Rowe  
Carnegie Mellon University  
Pittsburgh, PA  
agr@ece.cmu.edu

## ABSTRACT

Losses stemming from energy theft and system faults are a major challenge to providing reliable electrical service in developing areas of the world. Managing these losses is a vital part of ensuring energy distribution system stability and maintaining a functioning microgrid. Despite this, even losses that are detected are often not addressed by microgrid operators due to the difficulty of locating unauthorized loads, especially in deployments serving hundreds or thousands of households. In this paper, we propose a method for the localization of unauthorized loads based on the fine-grained sensing of voltage phase across a microgrid. Unlike other approaches, the proposed method utilizes synchronous voltage sensing at smart meters and does not rely on expensive inline power metering. Voltage phase measurements feed a graphical model of a power distribution network, which yields the locations of loads as they connect to the system. We evaluate our method using a circuit-based approach in SPICE by simulating loads on a real-world microgrid topology. We then validate our simulation results on a laboratory microgrid testbed using real loads, showing that fine-grained voltage sensing can be effectively leveraged to localize unauthorized loads in microgrids.

## CCS CONCEPTS

• Computer systems organization → Embedded and cyber-physical systems; • Hardware → Energy metering; Smart grid;

## KEYWORDS

Microgrid, Non-Technical Loss, Load Localization

## ACM Reference format:

Maxim Buevich and Anthony Rowe. 2019. Localizing Loads in Microgrids Using High-Precision Voltage Phase. In *Proceedings of 10th ACM/IEEE International Conference on Cyber-Physical Systems (with CPS-IoT Week 2019)*, Montreal, QC, Canada, April 16–18, 2019 (ICCPs '19), 10 pages. <https://doi.org/10.1145/3302509.3311037>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ICCPs '19, April 16–18, 2019, Montreal, QC, Canada

© 2019 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-6285-6/19/04...\$15.00

<https://doi.org/10.1145/3302509.3311037>

## 1 INTRODUCTION

Smart grid technologies in development today represent a real opportunity to revolutionize electrical power transmission and distribution in terms of efficiency, reliability, and sustainability. Innovations in digital communications and control have already led to gains by enabling the dynamic optimization of grid operations and the integration of distributed resources and generation, including renewables such as wind and solar power. The development and proliferation of a new generation of digital energy meters and connected appliances promise further gains through the use of demand-side management technologies while also enabling greater support for bi-directional energy flows and distributed resources at a low-voltage consumer level. However, smart grid architectures generally require connectivity to a national or international wide-area power grid, and thus are not suited for serving the more than 1.1 billion people in rural and remote areas without access to a reliable source of electricity [4]. Efforts to address this problem have introduced a new class of power distribution system called the "microgrid," a small-scale power grid with localized sources and loads. Microgrids are uniquely suited for serving geographically isolated areas because they are designed to operate both with or without a connection to a traditional wide-area grid. This key feature of microgrids presents new system requirements and challenges in addition to those of traditional power distribution.

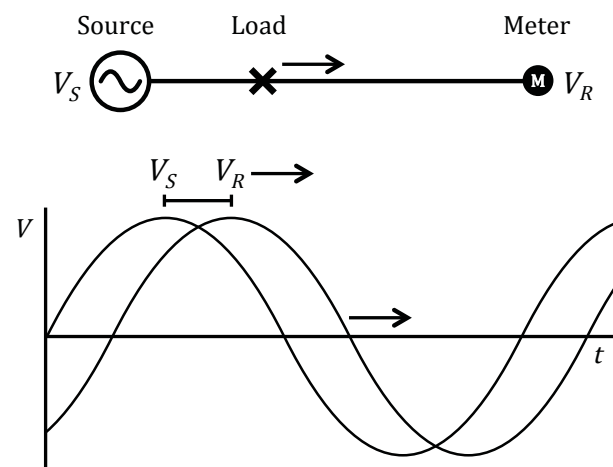


Figure 1: voltage phase shift across the length of a distribution line is related to the length of the current-carrying segment of the line

A major challenge of operating microgrids in developing areas is system reliability, which is made difficult by the pervasiveness of energy theft in these areas. Energy theft is often carried out by tampering with electrical meters or making unauthorized connections directly to the bare distribution wiring on utility poles. Such schemes bypass standard smart meters, making unauthorized connections difficult to locate. Due to the scale of some microgrids, even losses that are detected often go unaddressed because the locations of unauthorized loads cannot be determined.

In this paper, we propose a method for the localization of unauthorized loads based on the fine-grained sensing of voltage phase across a microgrid. Our method's use of voltage phase is a key feature in that it can be integrated into existing smart meters, and in that it avoids relying on cumbersome and expensive current sensing or inline power metering. Our method leverages the relationship between the voltage phase shift across the length of a distribution line and the length of the current-carrying segment of the line, as depicted in Figure 1. We propose a graphical microgrid model and discuss its derivation from a well-established model of transmission lines. We then present an algorithm which, in combination with the microgrid model and the nearest neighbor search (NNS) method, achieves accurate load localization. We evaluate the localization accuracy of our approach in a custom SPICE-based simulator using a subsection of a 500-household microgrid topology from Les Anglais, Haiti [5]. We analyze the effect of time synchronization on localization accuracy and show that a cost-effective wireless mesh network of smart meters can outperform a costly GPS-connected network of micro-phasor measurement units (micro-PMUs). Finally, we experimentally validate our graphical microgrid model and simulation results on a laboratory microgrid testbed using real loads.

## 2 RELATED WORK

The work presented in this paper utilizes concepts across multiple areas of systems research. In this section, we discuss related work in the areas of microgrid reliability, energy theft detection, power distribution fault localization, and power system phase measurement.

### 2.1 Microgrid Reliability

While microgrids can be an effective means of providing electrical power to areas outside the reach of a wide-area grid, they suffer from significant reliability and sustainability challenges. Research in the area of microgrid modeling has identified approaches for improving microgrid performance, including the use of renewable energy [13] [16] [12] and demand-side management strategies [11] [15]. The use of pre-paid smart meters has further enabled microgrid operators to effectively manage loads and collect tariffs [2] [1] [3]. Features such as remote system monitoring, load-shedding, and dynamic power-limiting allow operators to take reactive steps in cases of brown-out or system fault [11].

### 2.2 Energy Theft Detection

A number of approaches have been proposed for the detection of energy theft in microgrids and traditional power distribution systems. [18] provides an overview and analysis of common solutions, including the use of tamper-resistant meters, the tracking of

consumption over time, and line inspection. A proposed approach for preventing energy theft involves the injection of destructive harmonics into microgrid lines in order to damage unauthorized loads [8]. Some work has attempted to apply machine learning methods to smart meter data in order to perform theft detection. Support vector machines (SVM) have been employed to identify abnormalities in historical data but are limited by the fact that energy theft traces exhibit similar patterns to legitimate traces because they often involve the same types of appliances [7] [14]. In addition to existing limitations, most detection methods cannot account for the bypassing of energy meters and thus cannot detect the common case of a direct connection being made to the distribution wiring.

### 2.3 Power Distribution Fault Localization

Investigations into the localization of disturbances in power systems have been dominated by work at the high-voltage transmission level, as opposed to the low-voltage distribution level analogous to microgrids. Despite this, some methods for fault localization in unbalanced distribution systems have been explored. [9] utilizes the detection of high-frequency components caused by discharges from the low-level breakdown of insulators. This method is specific to electrical faults and would not be useful for localizing unauthorized loads. [10] presents an approach based on the estimation of apparent impedance, using fundamental components of voltage and current from a single location. This method is limited in that it requires current data and only provides a distance from the measurement point, as opposed to a location. [6] extends the method by utilizing multiple measurement points, but still requires current measurement. While all of the methods discussed were evaluated in simulation, none were tested experimentally.

### 2.4 Power System Phase Measurement

The work in this paper employs high-precision measurement of voltage phase across a microgrid in order to localize unauthorized loads. The measurement of phase in power systems is not a new concept, and is generally achieved through the use of phasor measurement units (PMUs). A PMU is a specialized device for measuring AC waveforms at a particular point in an electrical grid. A PMU captures the magnitude and phase of sinusoidal voltages and currents and represents them as phasors. A deployment generally consists of a network of many PMUs distributed across a grid, with all PMUs using a common time source for synchronization. Such a network of PMUs generates a set of time-synchronized phasor measurements known as a *synchrophasor*, which has become an invaluable tool for power management in electrical grids. Today, PMU networks are almost exclusively used to monitor high-voltage transmission systems, rather than mid- and low-voltage distribution systems. The latter are the focus of this work. Because both generation and loads are highly distributed at the transmission level, grid operators must rely on derived estimates of power flows in the system. These are generally found by measuring state variables throughout the grid using a PMU network and performing power-flow analysis, modeling the grid as a nonlinear system.

Historically, a couple of factors have limited the widespread adoption of PMU technology. High costs associated with installation, communication, and hardware have limited the size of PMU

networks at the transmission level. The 2014 report [20] by the U.S. Department of Energy found that the average overall cost per PMU unit (including procurement, installation, and commissioning) ranged from \$40,000 to \$180,000. Additionally, traditional PMU technology has seen almost no adoption at the distribution level on a local scale; this is due to both cost and the differing design requirements of distribution-level waveform sensing. For instance, a traditional PMU measures voltage phase with one-degree resolution, which works well for most transmission-level scenarios but is not useful at a local distribution scale, where much shorter lines carrying far less power result in much smaller voltage phase differences.

One initiative working to bring AC waveform measurement to distribution systems is the micro-PMU project. As discussed in [19], the micro-PMU project aims to develop a PMU architecture cost-effective and precise enough for use in distribution systems. A key contribution is the design of the micro-PMU hardware, which synchronizes measurements using GPS as a common time reference, achieving a voltage phase measurement resolution of 10 milli-degrees, equivalent to 0.46 microseconds at 60 Hz. In [19], the authors briefly discuss fault localization based on impedance estimates. However, this is presented as a hypothetical application and is not developed on or evaluated. While micro-PMU's precision of measurement is enough to localize unauthorized loads, the micro-PMU instruments presented are still prohibitively expensive for use in rural microgrids, with each unit costing approximately \$3,500. A deployment would involve additional costs likely exceeding those of the instrument, including costs incurred from unit installation, GPS antenna wiring, and cellular data usage. In contrast, this paper presents an unauthorized load localization method that is based on wireless energy meters costing less than \$100 per unit and does not require any additional instrumentation, GPS, or cellular connectivity.

### 3 GRID MODEL

Our load localization method exploits the fact that a load connected to a distribution line induces a voltage phase shift across the length of the line. In this section, we first describe the line model used as the basis for the relationship between load location and voltage phase shift; we then describe our model for graphically representing the electrical topology of a grid.

#### 3.1 Line Model

When a load draws power from a source through a transmission or distribution line, a certain portion of the total power - termed "line losses" - is consumed by the line itself. This phenomenon generally results in shifts of both the voltage amplitude and voltage phase throughout the length of the line. These shifts are determined by the properties of the load, the properties of the line, and the length of the line segment between the source and the load. As depicted in Figure 1, our method of load localization specifically exploits the relationship between the length of the line segment and shifts in voltage phase.

For the sake of simplicity and generalizability, we adopt a line model that assumes a linear relationship between the length of a current-carrying line segment and voltage phase shift. That is,

for every doubling of the line segment we expect a doubling of the phase shift. In reality, this is not the case and the relationship is in fact sub-linear for real distribution lines. The behavior does, however, closely approximate linearity for sub-kilowatt loads on low-voltage lines, leading us to adopt the simplified model. A standard model often used in power transmission literature that accurately describes the sub-linear relationship is the *short transmission line model*. In this section, we derive the expression for our linear model and show how its description follows from the *short transmission line model*. Later in this chapter, we explain how we use the *short transmission line model* as our "ground truth" model in SPICE simulation in order to evaluate the linear model.

**3.1.1 Relationship Between Voltage Phase and Load/Line Properties.** In order to describe, and later simulate, the sub-linear behavior of a real line, we employ the widely-adopted "short transmission line model." The model, shown in Figure 2, is a lumped-element circuit model for characterizing low-voltage transmission lines shorter than 80 kilometers. It is distinct from other transmission line models in that it ignores shunt capacitance, which tends to be relatively negligible for short lines. The short transmission line model enables us to calculate the load phasor voltage  $V_R$ , given a source phasor voltage  $V_S$ , a complex load impedance  $Z_R$ , and a complex line impedance  $Z_T = R_T + jX_T$ , where  $R_T$  is the line resistance and  $X_T$  is the line reactance. If we treat the circuit as an AC voltage divider, we find that

$$V_R = V_S \frac{Z_R}{Z_T + Z_R},$$

which can be represented in polar form as

$$|V_R|e^{i\omega t} e^{i\theta_{V_R}} = |V_S|e^{i\omega t} e^{i\theta_{V_S}} \frac{|Z_R|e^{i\phi_R}}{|Z_T|e^{i\phi_T} + |Z_R|e^{i\phi_R}}.$$

Our goal is to solve for the phase difference between the source voltage  $V_S$  and the load voltage  $V_R$ , as depicted in Figure 1. We do this by defining the source voltage phase to be zero,  $\theta_{V_S} = 0$ , and declaring a new impedance  $Z_F = Z_T + Z_R$ . We can then simplify the polar form




Figure 2: "short transmission line" model

$$\begin{aligned}
 |V_R|e^{i\omega t}e^{i\theta_{V_R}} &= |V_S|e^{i\omega t}e^{i0}\frac{|Z_R|e^{i\phi_R}}{|Z_F|e^{i\phi_F}} \\
 &= \frac{|V_S||Z_R|}{|Z_F|}e^{i\omega t}\frac{e^{i\phi_R}}{e^{i\phi_F}} \\
 &= \frac{|V_S||Z_R|}{|Z_F|}e^{i\omega t}e^{i(\phi_R-\phi_F)}
 \end{aligned}$$

to determine that

$$|V_R| = \frac{|V_S||Z_R|}{|Z_F|} \quad \text{and} \quad \theta_{V_R} = \phi_R - \phi_F.$$

The expression above shows that the load voltage phase  $\theta_{V_R}$ , and hence the voltage phase difference across a short transmission line, is solely dependent on the load phase shift  $\phi_R$  and the phase shift caused by the summation of load impedance and line impedance,  $\phi_F$ . A visual representation of this relationship can be seen in Figure 3, which shows "phasor diagrams" of how impedances determine load voltage phase  $\theta_{V_R}$  when the load is inductive (top) and when the load is capacitive (bottom).

Figure 3 also shows how the load voltage phase  $\theta_{V_R}$  is partially determined by the magnitude of the line impedance  $|Z_T|$ . This relationship is key to our localization method because  $|Z_T|$  is determined by the length of the line. This relationship is linear, thus when the length of the line is doubled the value of  $|Z_T|$  is doubled. There is one situation in which  $|Z_T|$  has no impact on the value of  $\theta_{V_R}$ : when the load impedance phase shift  $\phi_R$  and the line impedance phase shift  $\phi_T$  have the same value, then  $\theta_{V_R} = 0$  and the value of  $|Z_T|$  is irrelevant. In Figure 3, this would be visually represented by  $Z_R$  and  $Z_T$  pointing in the same direction. In practice, however, this case is unlikely to occur because a majority of real loads have a higher phase shift or lower phase shift than a low-voltage distribution line; that is, they have a considerably higher inductance per unit resistance than a distribution line (such as with




Figure 3: "phasor diagram" of complex impedances as they relate to voltage phase across a line (note: complex impedances are technically not phasors)

induction motors), or they have a considerably lower inductance per unit resistance (such as with incandescent bulbs and hot plates).

**3.1.2 Approximate Linearity of Voltage Phase.** For the purposes of load localization, we assume that when a load is powered through a distribution line, the length of a current-carrying line segment is linear and proportional to the voltage phase difference across that line segment. In reality, this relationship is sub-linear, but is still a close enough approximation of linearity for the purposes of our localization method. We can show that the relationship is approximately linear by modifying the polar form of our original definition of  $V_R$ , once again defining that  $\theta_{V_S} = 0$ ,

$$\begin{aligned}
 |V_R|e^{i\omega t}e^{i\theta_{V_R}} &= |V_S|e^{i\omega t}e^{i0}\frac{|Z_R|e^{i\phi_R}}{|Z_T|e^{i\phi_T} + |Z_R|e^{i\phi_R}} \\
 &= |V_S|e^{i\omega t}\frac{|Z_R|e^{i\phi_R}}{|Z_T|e^{i\phi_T} + |Z_R|e^{i\phi_R}} * \frac{e^{-i\phi_R}}{e^{-i\phi_R}} \\
 &= |V_S|e^{i\omega t}\frac{|Z_R|}{|Z_T|e^{i(\phi_T-\phi_R)} + |Z_R|},
 \end{aligned}$$

and declaring a new impedance  $Z_G$  to be

$$|Z_G|e^{i\phi_G} = |Z_T|e^{i(\phi_T-\phi_R)} + |Z_R|,$$

we simplify to

$$|V_R|e^{i\omega t}e^{i\theta_{V_R}} = |V_S|e^{i\omega t}\frac{|Z_R|}{|Z_G|e^{i\phi_G}}$$

and arrive at the definitions

$$|V_R| = \frac{|V_S||Z_R|}{|Z_G|} \quad \text{and} \quad \theta_{V_R} = -\phi_G.$$

Using the phasor arithmetic definitions for the phase angle of the sum of two phasors,

$$\phi_{A+B} = \arctan\left(\frac{|A|\sin(\phi_A) + |B|\sin(\phi_B)}{|A|\cos(\phi_A) + |B|\cos(\phi_B)}\right),$$

we can redefine  $\theta_{V_R}$  in terms of  $Z_T$  and  $Z_R$ ,

$$\begin{aligned}
 \theta_{V_R} &= -\phi_G \\
 &= -\arctan\left(\frac{|Z_T|\sin(\phi_T - \phi_R) + |Z_R|\sin(0)}{|Z_T|\cos(\phi_T - \phi_R) + |Z_R|\cos(0)}\right) \\
 &= -\arctan\left(\frac{|Z_T|\sin(\phi_T - \phi_R)}{|Z_T|\cos(\phi_T - \phi_R) + |Z_R|}\right)
 \end{aligned}$$

We further simplify the definition of  $\theta_{V_R}$  by noting that  $|Z_R| \gg |Z_T|$  is always true in a distribution line scenario. This is because loads powered through a line must have an impedance at least 20 times higher than that of the line (normally much higher), because otherwise the line cannot provide the rated voltage that is expected by the loads. This results in  $|Z_R|$  being the dominant term in the denominator, yielding the approximation

$$\theta_{V_R} \approx -\arctan\left(\frac{|Z_T|\sin(\phi_T - \phi_R)}{|Z_R|}\right).$$

This definition of  $\theta_{V_R}$  reveals the approximate linearity between the length of a current-carrying line segment and the voltage phase

difference across that line segment. Linearly scaling the length of a line segment causes  $|Z_T|$  to be scaled linearly, leaving all other parameters unchanged; this, in turn, causes the entire term within the  $\arctan()$  to be scaled linearly. Additionally, because  $|Z_R| \gg |Z_T|$ , it is known that the term within the  $\arctan()$  will equal a value near 0. Finally, the  $\arctan()$  function is known to approximate the linear function  $f(x) = x$  for values near 0, therefore we know that  $\theta_{V_R}$  will have an approximately linear response with respect to the length of a line.

### 3.2 Grid Graph Model

The load localization method described in this work is based on a comparison of voltage phase measurements collected from meters throughout the distribution grid. The method requires two sets of inputs: the first is the set of voltage phase measurements, and the second is a representation of the electrical topology (or "wiring") of the microgrid. We model the electrical topology as a directed rooted tree (or "in-tree"), where all edges point from a node to its parent and together comprise paths which all point to the root of the tree. As depicted in Figure 4, each node  $x_i$  in the tree represents a meter or a branching point in the grid, and each edge represents a line segment that connects a meter or branch point to a neighboring meter or branch point. The root node  $x_1$  represents the meter or branch point which is the immediate neighbor of the source in the grid. Each node  $x_i$  has two properties: a binary property  $m_i$  that is true when a node is a meter and false when the node is a branching point, and a numerical property  $d_i$  representing the line distance from the node to the source. The line distance is calculated from the physical length of the distribution lines when all line segments consist of the same type of wiring. However, the line distance can alternatively represent an impedance in cases where not all line segments have the same gauge or other electrical properties.

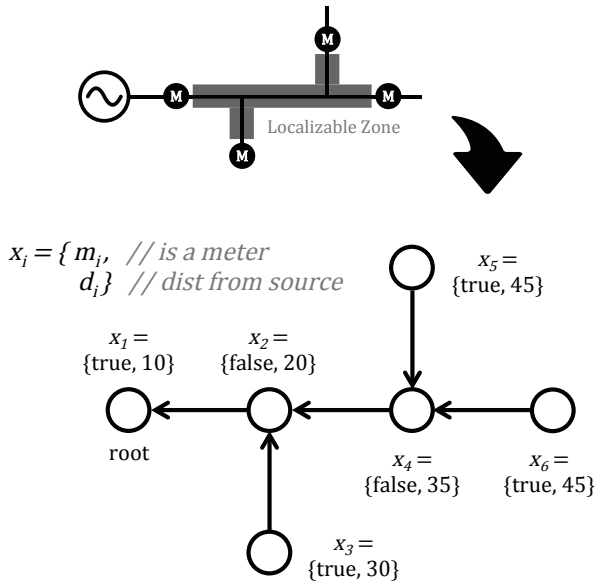


Figure 4: Example instance of the grid model generated from the electrical topology of a microgrid

The topological tree model is constructed from a subsection of the grid which we call the "localizable zone" of the grid. This zone represents the area of the grid in which our method can localize a load, and it is defined as the union of all paths between the location of the root node and the location of each individual meter in the grid. The set of line segments not in this zone consists of "leaf" segments which do not have a meter at the point of termination, and the "trunk" segment which connects the source to the root node. Any load present outside of the zone will be localized to the point in the zone closest to the real location. Grid developers can maximize the size of the localizable zone by ensuring that a meter is placed near the source and that all leaf line segments have a meter at the point of termination.

### 4 LOCALIZATION METHOD

**Algorithm 1:** The  $\text{Ph}()$  function generates node phases given a load location

```

Input: Finite set  $X = \{x_1, x_2, \dots, x_n\}$  of nodes,
         the node  $x_k$  associated with the load,
         the distance  $c$  of the load from the source
Output: Finite set  $P = \{p_1, p_2, \dots, p_n\}$  of normalized node
         voltage phases
1 declare empty set  $P$ 
2 for  $i \leftarrow 1$  to  $n$  do
3    $\text{ancestor} \leftarrow \text{lowest\_common\_ancestor}(x_i, x_k)$ 
4   if  $\text{ancestor} = x_k$  then
5      $p_i \leftarrow c$ 
6   else
7      $p_i \leftarrow \text{distance}(\text{ancestor})$ 
8 for  $i \leftarrow 1$  to  $n$  do
9    $p_i \leftarrow p_i / \text{largest\_magnitude\_in}(P)$ 
10 return  $P$ 

```

In order to localize loads in the grid, we first define a phase function  $\text{Ph}()$  which, given a load location, calculates the expected

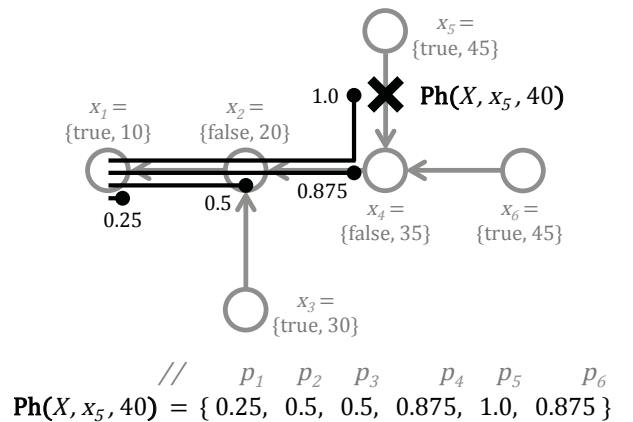


Figure 5: Example of the phase function output given an electrical topology and load location

relative voltage phases at each node (see Algorithm 1). As inputs, the function takes the set of nodes  $X$ , the node  $x_k$  associated with the load, and the line distance  $c$  between the load and the source. The node  $x_k$  associated with the load is defined as the child node of the line segment the load is connected to. For each node  $x_i$  in the grid, the expected voltage phase  $p_i$  at that node is determined by the *lowest common ancestor (LCA)* of the node  $x_i$  and the load-associated node  $x_k$ . The lowest common ancestor of two nodes is the deepest node which has both nodes as descendants. (Of note is the fact that if one of the nodes is a descendant of the other then that other node is the LCA. This is because each node is defined to be a descendant of itself.) Each of the voltage phases  $p_i$  is then scaled by the voltage phase with the largest magnitude. The result is the set  $P$  consisting of relative phases  $p_i$  with values between 0 and 1, which is returned by the function. Figure 5 shows a simple example of the output of  $\mathbf{Ph}()$  given a load connected 40 units away from the source.

Using the phase function  $\mathbf{Ph}()$ , we then construct a search space that can be queried using a host of optimization techniques. For the sake of simplicity, we employ the widely-adopted *nearest neighbor search (NNS)* technique in order to localize a load. The search space is represented by the finite set  $S$ , where each  $s_i$  is itself a set and equals the output of the phase function  $\mathbf{Ph}()$  at a given load location. The set  $S$  is constructed by starting with a load at the root node and linearly scanning the load across all line segments in the grid, using a chosen interval. For example, in order to perform localization at a 1-meter resolution one would set the interval to 1 meter. If the grid in this scenario consists of a total of 1000 meters, then the set  $S$  will contain 1001 sets of relative phases.

Once a search space  $S$  is defined, it can be used to localize a given query point  $q$ . When a load is connected to an unpopulated grid,  $q$  is the set of phase measurements from each meter in the grid, normalized such that they can be directly compared to the sets  $S$ . When a load is connected to a populated grid,  $q$  is the difference in phase at each node between the current measurement and the previous measurement, taken before the appearance of the load. To determine the location corresponding to  $q$ , we calculate the dissimilarity for each  $s_i$  and  $q$  using the Euclidean distance function

$$d(s'_i, q) = \sqrt{(s'_{i,1} - q_1)^2 + (s'_{i,2} - q_2)^2 + \dots + (s'_{i,n} - q_n)^2},$$

where  $s'_i$  is the subset of  $s_i$  containing phase estimates of meter nodes, and excluding those of branch nodes. The location corresponding to the  $s'_i$  with the lowest dissimilarity then indicates the location of  $q$ .

## 5 EVALUATION

In this section we describe our synthetic load model, evaluate our localization method at scale in simulation, and discuss experimental validation on an in-lab testbed.

### 5.1 Synthetic Load Model

While electrical devices and appliances are often thought of as drawing a stable amount of power over time, in reality a load's power consumption is constantly varying. These variations in power cause proportional variations in voltage phase across a grid. Figure 6

shows example traces of this phase jitter for (a) a 72-Watt incandescent light bulb and (b) a 50-Watt floor fan. These traces were collected with the loads powered through 30.48 m (100 ft) of 16 AWG two-conductor wire. The incandescent bulb is shown to exhibit a voltage phase that varies by approximately 5 nanoseconds (or 5% of the phase), while the voltage phase of the floor fan varies by a significantly larger 60 nanoseconds (10% of the phase). This voltage phase jitter is the dominant source of error for the localization method presented in this work; hence, understanding and accurately modeling this phenomenon is a critical part of our evaluation.

We build on top of the load model in the "short transmission line" case by representing loads as random impedance variables. This synthetic load model is not used by our localization method but rather is used by the simulator to mimic the "noisiness" of real loads. Hence, the model is presented as a part of evaluation. A synthetic load is defined as a random variable

$$D \sim Z * N(\mu, \sigma^2),$$


$$\text{where } Z = R + jX \text{ and } \mu = 1.$$

A synthetic load  $D$  consists of a complex impedance  $Z$  and a normally-distributed random noise function  $N(\mu, \sigma^2)$  with mean  $\mu$  equal to 1 and variance  $\sigma^2$  determining the amount of jitter. We have found that for both linear and non-linear loads, the measured voltage phase tends to be normally distributed and closely matches our load model.


**Table 1: Load Set**

Load	$R$ ( $\Omega$ )	$X$ ( $\Omega$ )	$\sigma$
Incandescent Bulb	201.0	0.0	0.009
Floor Fan	288.9	22.3	0.016
Desk Fan	558.1	98.2	0.013
LED Bulb	704.2	54.1	0.032

**5.1.1 Load Set.** In order to populate our microgrid simulator, we characterized a small collection of loads we call the *Load Set* using a laboratory testbed. Devices were powered through 30.48 meters (100 feet) of two-conductor 16 AWG wire and high-precision voltage measurements were collected in order to estimate load impedance and noise of each device. The *Load Set* is shown in Table 1 and consists of a 72-Watt incandescent light bulb, a 50-Watt floor fan, a 25-Watt desk fan, and a 20-Watt LED bulb. These particular loads were chosen to represent the three major categories of loads: purely-resistive linear loads (incandescent bulb), inductive/capacitive linear loads (fans), and non-linear loads (LED bulb). These loads were also chosen because they are commonly used by customers in rural microgrids. Each load in the load set is characterized by a resistance  $R$ , a reactance  $X$ , and a variance  $\sigma$ . The properties of the loads in the set are used by the simulator to generate synthetic loads for localizing and for simulating background noise. One special case is the LED bulb, which is a non-linear load that we are representing using its linear properties. While doing



(a) 72-Watt incandescent light bulb



(b) 50-Watt floor fan

**Figure 6: Measured voltage phase jitter distribution for incandescent bulb and floor fan**

this does not yield an accurate representation of the LED bulb voltage waveform in simulation, we found through experimentation on our testbed that we still achieved accurate voltage phase measurements through zero-crossing detection. We believe that the handling of arbitrary non-linear loads would require a redesign of the microgrid simulator and is a promising direction for future work on this topic.

Employing different load types, we found that both resistance and reactance (due to capacitance or inductance) have a significant effect on voltage phase. This can be seen in Figure 6, which shows that the voltage phase shift and phase noise of (b) the 50-Watt floor fan is significantly higher than that of the (a) 72-Watt incandescent bulb, despite the fact that the bulb draws more power. This illustrates a counterintuitive result where the localization method based on voltage phase will localize a low-power load more accurately than a high-power load in certain cases. One such corner case is illustrated, in which there is only a single load connected on the microgrid. Note that both Table 1 and Figure 6 are calculated using phase data that has some filtering already applied; specifically, each phase reading is calculated based on the average of 60 zero-crossings over a 1-second period, where the zero-crossings are passed through an RC low-pass filter with a cut-off frequency of 11.4 kHz.

## 5.2 Simulator Design

In order to evaluate load localization at scale, we developed a custom SPICE-based microgrid simulator for simulating arbitrary grid topologies. The simulator consists of two stages: the SPICE stage and the localization stage.


The inputs to the simulator are fed into the SPICE stage, consisting of a graphical model of the microgrid electrical topology (or "wiring") and the *Load Set*. The topological information is used to build a circuit representation of the grid, where each 1-meter segment of wire is modeled as a "short transmission line" component. The electrical topology also defines the locations of energy meters in the grid. Loads are represented by the synthetic load model, through which they are described by a characteristic resistance  $R$ , reactance  $X$ , and standard deviation of noise  $\sigma$ . The user chooses a target load and a set of background loads from the *Load Set*, as well as the locations of the loads. The set of loads and locations can be either pre-defined or chosen randomly by the simulator. The full circuit is then built and the steady-state AC response is simulated, generating the set  $q$  of voltage phases at energy meter locations.




**Figure 7: Electrical topology of Les Anglais microgrid subdivision**

The output of the SPICE stage is used as the input to the localization stage of the simulator. As described in Section 4, this stage employs the topology input and the phase function  $\mathbf{Ph}()$  to construct a search space  $S$ . The simulator then linearly searches the set within  $S$  that has the lowest dissimilarity to the query set  $q$ . Although linear search of a space is often considered inefficient, we have found that it works well in this case because the search space is small, growing proportionally with the sum of the lengths of the line segments in the grid.

For our evaluation, we used a subdivision of the microgrid in Les Anglais, Haiti, shown in Figure 7, as the electrical topology input of our simulations. The entire microgrid consists of five single-phase low-voltage subdivisions powered by the three-phase medium-voltage distribution system through five separate transformers. Due to voltage distortions across transformers, the voltage phase-based method presented in this work cannot localize loads across transformers, requiring that each subdivision be handled as a separate microgrid. For the sake of practicality, a single subdivision was employed in simulation. The subdivision consists of 800 meters of two-conductor 2 AWG wire, separated by approximately 2 meters, carrying single-phase 120-Volt AC power. In the physical microgrid, this low-voltage line is strung along utility poles, represented by the white line segments and numbers in Figure 7. The location of the source is indicated by a large white point and the locations of the meters are indicated by small red points.



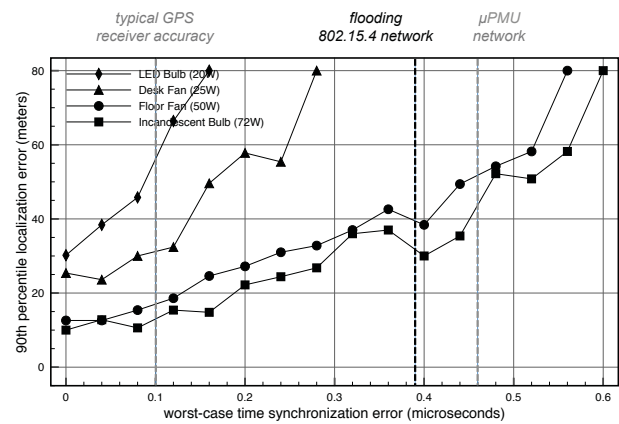
(a) 10 random background loads



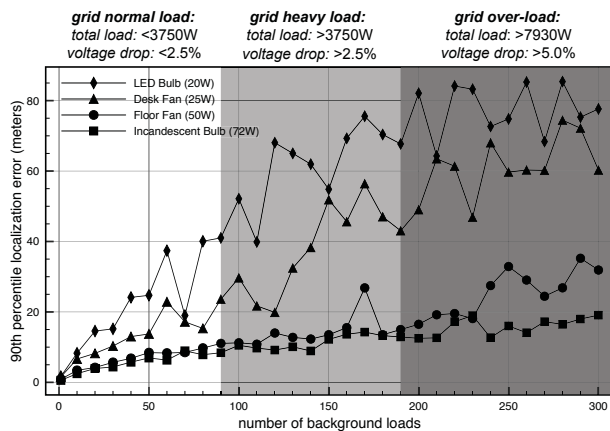
(b) 100 random background loads

**Figure 8: Location estimate distributions for a 72-Watt incandescent light bulb given 10 background loads and 100 background loads**

**5.2.1 Simulation Results.** Through simulations on the synthetic Les Anglais electrical topology, we aimed to capture the relationship between localization error and background noise on the lines, and provide insight into how well we could expect to localize loads on a real microgrid at run-time. Figure 8 shows a comparison of localization error between a low and high background noise scenario. Figure 8(a) represents the distribution of localization error for the 72-Watt incandescent bulb connected 100 meters from the source on a grid with 10 random background loads; Figure 8(b) shows the same situation with 100 background loads. We see that the error distribution is significantly larger when there is more background noise due to a larger number of loads. We also see that in both cases in Figure 8, our localization method tends to overestimate the distance of the load. This is a general behavior of our method that applies to all cases of localization, and is a result of our model’s assumption that load distance and phase have a linear relationship, which is in fact an approximation of the real relationship. This bias is more significant under heavy load and could potentially be accounted for with a more sophisticated model.



**Figure 10: 90th percentile localization error (in meters) with respect to actual load location, for different load types given the quality of time synchronization of devices. (Lower localization error is better. More accurate time synchronization results in more accurate load localization.)**



**Figure 9: 90th percentile localization error (in meters) with respect to actual load location, for different load types given a variable number of background loads. (Lower localization error is better. Increasing the number of background loads results in less accurate load localization.)**

We also investigated how different load types compare in terms of localization error, given an increasing amount of background noise. This was done by simulating each type of load on the grid topology with a gradually increasing number of background loads, from 1 to 300. The target load was placed 100 meters from the source, while the background loads were randomly chosen from the *Load Set* and placed at random locations. The results in Figure 9 show how 90th percentile localization error increases for each load type as the number of background nodes increases (90% of errors were below the given threshold). Each data point represents the 90th percentile error out of 100 simulations. We see that while localization error for all loads tends to grow logarithmically with the number of background nodes, loads vary greatly in terms of absolute localizability. The results suggest that load power is the main factor determining localization error, with the 72-Watt incandescent bulb exhibiting an error of up to 8.92 meters with the grid under normal load, while the 20-Watt LED bulb exhibits an error up to 40.07 meters. Errors for all load types are higher under heavy load, when the voltage drops by more than 2.5%, and higher still for over-load,



when the voltage drops by more than 5%. In practice, the overload case should be very uncommon because 120-Volt at 60 Hz systems are generally specified to voltage tolerances of no more than  $\pm 5\%$ . These results would be different for systems such as those conforming to the European Union's 230 Volt  $\pm 10\%$  at 50 Hz standards.

Finally, we investigated the impact of time synchronization across the metering network on localization error. This was achieved by adding a module to the simulator that would add a specified amount of random jitter to each phase value that was calculated during the SPICE simulation stage. The module models timing jitter as a Gaussian random variable and applies it one of two ways, using the "broadcast" mode or the "multi-hop" mode. The "broadcast" mode involves applying jitter equally to every node and was used to represent time synchronization in a network of GPS-connected micro-PMUs, while the "multi-hop" mode represents time synchronization in a mesh network, where synchronization of a node depends on the number of hops between the node and the gateway.

Figure 10 shows how the 90th percentile localization error for each load type exhibits an exponential-like increase as the time synchronization worsens, with the timing jitter being applied using the "broadcast" mode. We define quality of time synchronization in terms of worst-case synchronization error, where the worst-case error is equal to three standard deviations from the correct time. Loads were simulated on the microgrid subsection, 100 meters from the source, and with 100 random background loads placed at random locations. In both Figure 9 and Figure 10, variations in the results are due to each data point being derived from the 90th percentile of 100 simulations, requiring a significant amount of time to execute. "Smoother"-looking results can be generated by running more iterations of the simulation.

Dashed grey lines indicate the localization performance assuming the timing resolution of a micro-PMU network and a "best-case" scenario in which raw GPS receiver timing is guaranteed. The dashed black line indicates the localization performance and the equivalent "broadcast" worst-case time synchronization error of our multi-hop 802.15.4 network, equal to 390 nanoseconds for the topology used in these simulations. Equivalent "broadcast" synchronization performance was determined by simulating timing jitter in the "multi-hop" mode and matching the two modes in terms of equivalent localization error. Our multi-hop 802.15.4 network, characterized by a worst-case time synchronization error of 240 nanoseconds per hop, was found to have an equivalent "broadcast" mode worst-case error of 390 nanoseconds [17]. Figure 10 shows that this enables the network to localize the 72-Watt bulb and the 50-Watt fan to approximately 30 and 40 meters, respectively. While our network is not as accurate as the "best-case" GPS receiver network, which achieves time synchronization within 100 nanoseconds, it does outperform the micro-PMU network, which has a worst-case error of 460 nanoseconds and is considerably more costly.

### 5.3 Experimental Validation

We experimentally validated our models and simulation results using a small, in-lab microgrid testbed, shown in Figure 11. The testbed consists of 5 segments of 30.48-meter (100 feet), two-conductor





Figure 11: Microgrid testbed diagram and photo

16 AWG wire connecting two metering points on two separate branches to an AC power source. Phase data at the source and two metering points was collected using a custom high-voltage comparator circuit designed to measure voltage zero-crossings and output a square wave signal. The square wave signal was sampled by a logic analyzer at 500 mega-samples per second, and the data was collected by a laptop computer.

Using the testbed, we were able to validate our line model and load model, including the fact that voltage phase and voltage phase error increase approximately linearly with the line distance of a load. In addition, we were able to test our localization method with real loads. Figure 12 shows distributions of localization results as multiple incandescent bulbs are connected to the testbed one-by-one. Each data point in Figure 12 represents a 1-second sample calculated from the average of 60 zero-crossings which have been filtered through a passive RC filter with a cutoff frequency of 11.4 kHz. The example also shows the effect of undesirable phase shift caused by temperature change in the measurement circuit components, which could be corrected for through calibration in future iterations of the hardware. Despite these effects, we found that overall testbed localization results were both accurate and precise enough to warrant future evaluation in a full-scale deployment.

## 6 CONCLUSIONS

In this paper, we presented a method for localizing unauthorized loads through the high-precision measurement of voltage phase. The use of voltage, rather than current, makes this a relatively

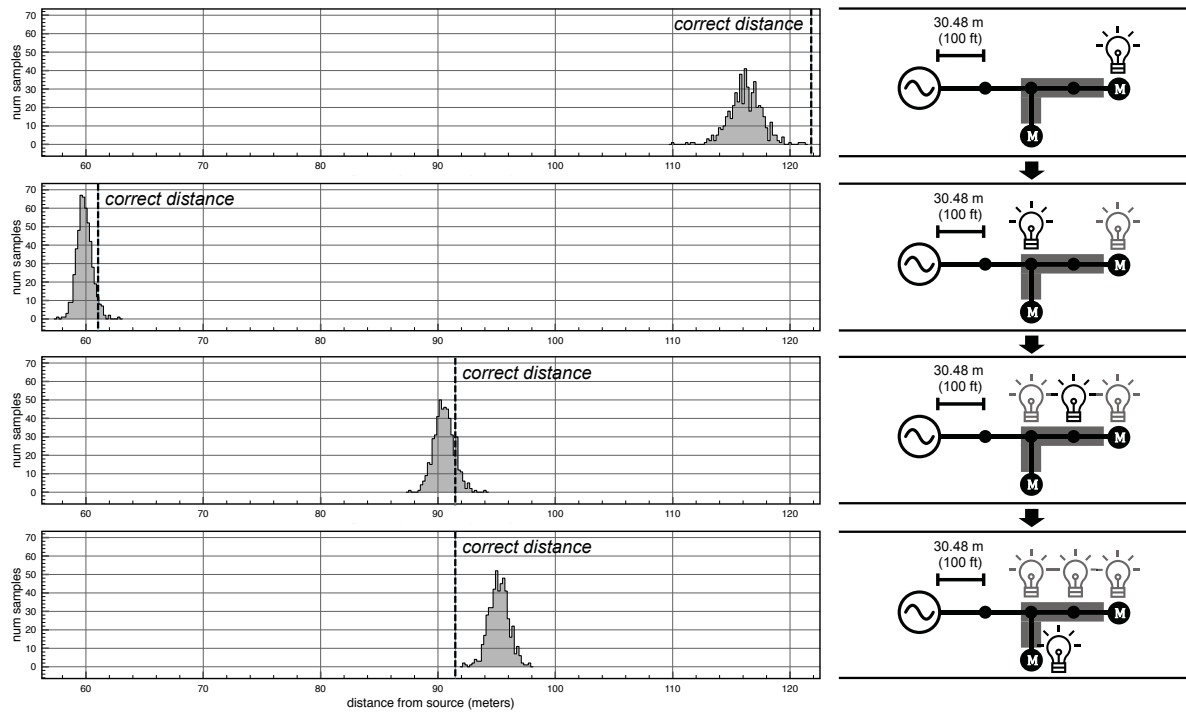


Figure 12: Localization error of multiple loads running on the microgrid testbed

low-cost and unintrusive method for use with smart meters in microgrids. We evaluated our approach using SPICE-based simulation on a real-world microgrid topology and validated our results on a laboratory microgrid testbed. Through our evaluation, we gained a deeper understanding of the relationship between background noise on lines, time synchronization, and localization error, and gained insight into the localization method’s potential real-world performance.

REFERENCES

[1] [n. d.]. Circutor 2012a. Renewable Energies: <http://circuitor.com> (viewed 2/22/2013). ([n. d.]).

[2] [n. d.]. Devergy: <http://www.devergy.com> (viewed 10/10/2013). ([n. d.]).

[3] [n. d.]. INENSUS: <http://www.inensus.com/en/products5.htm> (viewed 10/10/2013). ([n. d.]).

[4] International Energy Agency. 2017. World Energy Outlook database. (2017).

[5] Maxim Buevich, Xiao Zhang, Oliver Shih, Dan Schnitzer, Tristan Escalada, Arthur Jacquiau-Chamski, Jon Thacker, and Anthony Rowe. 2016. Microgrid losses: when the whole is greater than the sum of its parts. In *Proceedings of the 7th International Conference on Cyber-Physical Systems*. IEEE Press, 46.

[6] R Das, MS Sachdev, and TS Sidhu. 2000. A fault locator for radial subtransmission and distribution lines. In *Power Engineering Society Summer Meeting, 2000. IEEE*, Vol. 1. IEEE, 443–448.

[7] S.S.S.R. Depuru, Lingfeng Wang, and V. Devabhaktuni. 2011. Support vector machine based data classification for detection of electricity theft. In *Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES*. 1–8. <https://doi.org/10.1109/PSCE.2011.5772466>

[8] Soma Shekara Sreenadh Reddy Depuru. 2011. Electricity theft: Overview, issues, prevention and a smart meter based approach to control theft. *Energy Policy* (2011).

[9] M El-Hami, LL Lai, DJ Daruvala, and AT Johns. 1992. A new travelling-wave based scheme for fault detection on overhead power distribution feeders. *IEEE Transactions on Power Delivery* 7, 4 (1992), 1825–1833.

[10] Adly A Girgis, Christopher M Fallon, and David L Lubkeman. 1993. A fault location technique for rural distribution feeders. *IEEE Transactions on Industry Applications* 29, 6 (1993), 1170–1175.

[11] Meg Harper. 2013. Review of Strategies and Technologies for Demand-Side Management on Isolated Mini-Grids. (2013).

[12] Daniel M Kammen and Christian E Casillas. 2011. The Delivery of Low-Cost, Low-Carbon Rural Energy Services. (2011).

[13] M. Khan and M. Iqbal. 2005. Pre-feasibility study of stand-alone hybrid energy systems for applications in Newfoundland. *Renewable Energy* (2005).

[14] J. Nagi, K.S. Yap, S.K. Tiong, S.K. Ahmed, and A.M. Mohammad. 2008. Detection of abnormalities and electricity theft using genetic Support Vector Machines. In *TENCON 2008 - 2008 IEEE Region 10 Conference*. 1–6. <https://doi.org/10.1109/TENCON.2008.4766403>

[15] TG Quetchenbach, MJ Harper, J Robinson IV, KK Hervin, NA Chase, C Dorji, and AE Jacobson. 2013. The GridShare solution: a smart grid approach to improve service provision on a renewable energy mini-grid in Bhutan. *Environmental Research Letters* 8, 1 (2013), 014018.

[16] S Rana, R Chandra, SP Singh, and MS Sodha. 1998. Optimal mix of renewable energy resources to meet the electrical energy demand in villages of Madhya Pradesh. *Energy Conversion and Management* (1998).

[17] Thomas Schmid, Prabal Dutta, and Mani B Srivastava. 2010. High-resolution, low-power time synchronization an oxymoron no more. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*. ACM, 151–161.

[18] T. Smith. 2004. Electricity theft: a comparative analysis. *Energy Policy* (2004).

[19] Alexandra von Meier, Emma Stewart, Alex McEachern, Michael Andersen, and Laura Mehrmanesh. 2017. Precision micro-synchrophasors for distribution systems: A summary of applications. *IEEE Transactions on Smart Grid* 8, 6 (2017), 2926–2936.

[20] Marcus Young and Alison Silverstein. 2014. *Factors Affecting PMU Installation Costs*. Technical Report.