

# Grid Watch: Mapping Blackouts with Smart Phones

Noah Klugman<sup>†</sup>, Javier Rosa<sup>‡</sup>, Pat Pannuto<sup>†</sup>, Matthew Podolsky<sup>‡</sup>,  
William Huang<sup>†</sup>, and Prabal Dutta<sup>†</sup>

University of Michigan  
Ann Arbor, MI 48109

University of California, Berkeley  
Berkeley, CA 94720

{nklugman,ppannuto,wwhuang,prabal}@umich.edu {javirosa,podolsky}@eecs.berkeley.edu

## ABSTRACT

The power grid is one of humanity’s most significant engineering undertakings and is essential in developed and developing nations alike. Currently, transparency into the power grid relies on utilities and more fine-grained insight is provided by costly smart meter deployments. We claim that greater visibility into power grids can be provided in an inexpensive, crowdsourced manner independent of utilities by leveraging existing smartphones. Our key insight is that an unmodified smartphone can detect power outages by monitoring changes to its own power state, locally verifying these outages using a variety of sensors, and corroborating with other phones through cloud services. This approach enables a decentralized system that can scale, potentially providing researchers and concerned citizens with a powerful new tool to analyze the power grid and hold utilities accountable for degradation in power grid quality.

## Categories and Subject Descriptors

B.8.m [Hardware]: PERFORMANCE AND RELIABILITY—*Miscellaneous*

## General Terms

Economics, Measurement, Reliability, Security

## Keywords

Smart Grid, Power Monitoring, Crowdsourcing, Smartphone Applications, Side Channel Information

## 1. INTRODUCTION

The power grid is of enormous importance to global welfare, and it stands to reason that information regarding its stability would be of interest to researchers, policy makers, and the public at large.

The current paradigm for increasing visibility into the power grid is via a centralized network of utility owned, deployed, and controlled smart meters. While this can offer a highly detailed view into the power grid, this approach has its flaws. As the smart meters are controlled by utility companies, they do not necessarily yield greater transparency and visibility to researchers and the public at large, limiting their usefulness in helping third parties audit power grids. This is especially problematic in countries where corruption may play a role in controlling the external perception and reporting of power grid quality. Furthermore, smart grids are costly, and thus hard to scale in developing nations, which conversely is where power grids are least stable and where increased visibility may be the most useful.

We propose Grid Watch, a new bottom up, automated, and crowdsourced method of characterizing power grid stability. The key insight of Grid Watch is that smartphones can cheaply detect power outages by monitoring changes in state while charging. Furthermore, smartphones have various sensors which allow them to locally verify that a loss of power while charging is actually a power outage.

Grid Watch provides greater transparency as compared to smart meters by collecting data in a decentralized, grassroots manner, making it potentially more useful for monitoring, vetting, and auditing utilities. Furthermore, Grid Watch leverages the potent and still blossoming global smartphone community to cheaply scale, allowing it to occupy a niche in developing countries that may not be able to afford large smart meter deployments or to fill the void in countries such as the United States where the deployment of smart meters has been very slow.

To the best of our knowledge, there is no publicly available repository of power outage data, much less one that is automatically updated and independent from utility reports. We believe Grid Watch could play an integral role in the creation of such a data set, resulting in a generational shift in how third parties can analyze power grids and hold utilities accountable.

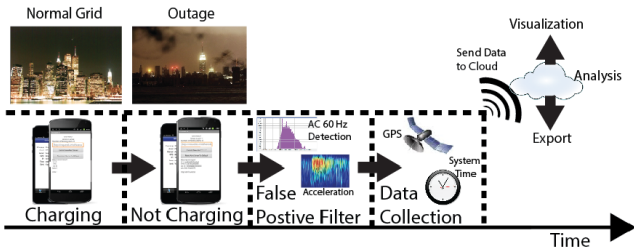


Figure 1: **Grid Watch Operation Diagram**— A plugged-in phone changes from a powered state to an unpowered state with grid failure. Grid Watch registers this event, verifies that it is not a false positive, and reports the event to the cloud for analysis, export and visualization.

## 2. GRID WATCH SYSTEM

“Grid Watch” is a crowd-sourced, automated, mobile sensing application. Grid Watch senses a power outage by taking advantage of two observations: 1) a phone is nearly never unplugged without being picked up and moved soon after, and 2) the “hum” of AC mains power can be detected using the microphone present on the phone. When Grid Watch detects that a phone has stopped charging, it samples from the accelerometer and microphone. If analysis of these samples show that an outage did occur, Grid Watch uploads the GPS location, system time, and phone unique ID to a central service. This data is prepared by the central service for export and visualization and is used as input data into grid behavior modeling algorithms. An overview of the system operation is shown in Figure 1. We implemented Grid Watch as a smartphone app for both Android and iOS.

### 2.1 Smartphone Power Outage Detection

Both Android and iOS expose charge state events which are used to wake the Grid Watch app up from the background. When an OS event registers that a phone has stopped charging, Grid Watch samples the accelerometer and microphone for 5 seconds. The accelerometer detects if the phone is being moved (unplugged), and an FFT on the microphone sample detects the AC mains hum. In addition, Android’s API exposes the classification of charger type, allowing Grid Watch to easily filter out charge state events that occur when the phone is charging in the car or over USB. Grid Watch reports the results of these tests to its central service. The app additionally allows users to manually report outages that were not automatically detected by Grid Watch (e.g. an outage that occurred when the user’s phone was not plugged it).

These same tests can be used in reverse to detect the power outage recovery rate, especially by infrequently sampling the microphone. Sampling five seconds of audio every 30 minutes would require only a 0.27% duty cycle while giving us currently unavailable data. There is demand for data on grid recovery, particularly in the event of natural disasters such as the aftermath of Hurricane Sandy [27].

### 2.2 The Data

Our current Grid Watch implementation collects the following data for each event, which we consider to be the bare minimum for Grid Watch to be effective:

**GPS Location:** In order to ascertain outage area, the location of outage events must be recorded. In deference to user privacy concerns, however, the GPS granularity is user controllable. While precise GPS data allows for high precision of outage reports, we hypothesize that a high density of low precision locations could also provide sufficiently accurate outage maps while preserving user privacy.

**Unique ID:** An ID is not strictly necessary for the correct operation of a Grid Watch. We collect it, however, for the purpose of estimating Grid Watch user base and density in a given area. In addition, we believe users may have an interest in tracking their own power outages. This belief is based on the prevalence of power outage maps available on utility websites in the United States. Lastly, this enables us to delete all of the events reported by a user if so requested.

**Classifier Results:** Both of our local outage filters are threshold-based (did the accelerometer move “too much”; is the magnitude of the 120 Hz peak “high enough” above the baseline?). We collect this baseline data to refine our classifiers and validate our thresholds.

**System Time:** When a potential outage event is detected, Grid Watch timestamps the power loss before taking any other action. This local timestamp is used as the Grid Watch ground truth of when an outage occurred. In addition, the Grid Watch service records a timestamp when the event is actually received. A survey of this delay provides a mechanism to test our hypothesis that the independent cellular backhaul remains a responsive means for reporting during outages.

### 2.3 The Central Service

Currently, the Grid Watch central service is responsible for archiving the reported outage data, basic access controls to the data, and providing users with feedback regarding their current power outage and power outage history.

## 3. EVALUATION

Although we have not yet deployed Grid Watch, we performed experiments to validate a subset of our system’s requirements. We use Grid Watch to detect a power outage in a house. In addition, we evaluate the Grid Watch false positive filters and measure the time synchronization between smartphones.

### 3.1 False Positive Detection

We performed two experiments to validate our methods of false positive detection. First, we ran an instance of Grid Watch using only the accelerometer to filter false positives on both an iPhone 5 for two weeks and a Galaxy Nexus smartphone for three days. During this time, phones were used routinely. The central service did not receive any false positives during this experiment.

Additionally, we tested the ability of phones to detect the 60 Hz hum from AC mains. We recorded a five second audio sample from four different phones inside a house and then turn off the master circuit breaker to simulate a power outage. We performed a basic FFT to search the signal for the AC mains frequency. The results of this survey are shown in

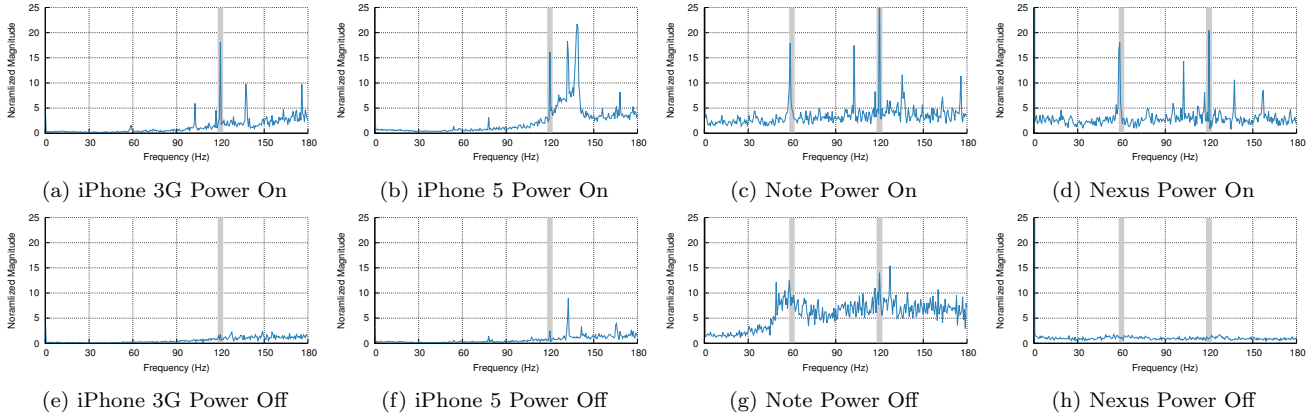



Figure 2: **FFT of audio samples captured on an array of phones inside a house**— The top row captures the environment during normal power and the bottom row is the same environment during a power outage. On all models, a 120 Hz peak (2nd harmonic of the 60 Hz mains) is visible when the power is on but disappears in the event of a power outage. The iPhones appear to have a 60 Hz notch filter installed in the audio frontend to mitigate noise from the environment, but none of the harmonics are filtered, allowing our microphone-based detection to remain effective.



(a) Ten phones detecting the same simulated power outages

Figure 3: **Exploring how tightly-coupled timing event detections are between a variety of phones**— For each trial, phones were connected to the same power outlet, which we switched off to simulate a power outage. Phones time stamped the event detection, and we characterize the standard deviation in reported times for each trial.

Figure 2. We observed the 60 Hz (U.S. AC mains’) frequency clearly on the phones running Android when the power in the house is on. On the iPhones, we observed no 60 Hz peak, indicating that a notch filter is present in their audio frontend, presumably because 60 Hz interference is undesirable for normal use of the microphone. Fortunately this filter does not extend to the second harmonic of the AC mains signal at 120 Hz. The 120 Hz harmonic is also highly detectable in on the Android phones. We built our AC mains classifier to detect the presence of a 120 Hz peak. To ensure accurate detection in other countries, our classifier also reports the presence of mains power if a 100 Hz peak, the second harmonic of a 50 Hz mains, is detected.

### 3.2 Time Synchronization Between Phones

Cascading outages spread within minutes [20], which indicates that the time resolution of Grid Watch outage reports must be sufficiently high fidelity to be useful in tracking the spatial and temporal spread of certain outages.

To test this, we connect ten smartphones, with six different models, to a single power strip. We turn off the power switch, and examine the time stamps each phone generates upon detecting the power outage. We repeated this process twenty times, and display the results in Figure 3.

We find that the average standard deviation between reported times for an identical power outage was .76 s, with the maximum standard deviation being 2.145 s. For 19 out of 20 trials, the standard deviation between reported times is less than .8 s. The maximum time difference between time stamps for a single event is 6.21 s. For 19 out of the 20 trials, the maximum time difference is less than 2.2 s.

From this, we conclude that the time synchronization between smartphones could be sufficient to help characterize power outage spreads. Our general knowledge of the dependence of GSM and GPS on accurate timing also supports the idea that the time synchronization of smartphones should be relatively high.

## 4. RESEARCH QUESTIONS

Our initial work in developing Grid Watch raises a number of research questions both for crowd sourced sensing systems at large and power for our application of grid monitoring. In this section we seek to enumerate what we see as the greatest challenges moving forward with Grid Watch and community sensing in general.

### 4.1 Grid Modeling

Power companies in developed countries use a combination of automated and manual techniques to identify and localize power outages. Battery-backed smart meters report outages, but their penetration is limited even in developed countries [18, 29]. As a result, customer provided reports supply the most actionable data but are neither automatic nor quick. We hypothesize that Grid Watch may be able to provide utility companies with information of high enough fidelity to support their efforts in performing demand response, energy consumption scheduling, and recovery from

massive power outages [20, 21]. We are uncertain, however, what penetration is necessary to provide results of sufficient quality to be actionable.

It would be desirable to be able to track the spread of an outage using Grid Watch data. Clustering appears to be a natural choice for this problem because the data is inherently clustered by geographic location and temporal position as well as by the topology of the power grid [23]. Additionally, past work has shown cascading power failures follow spatial and temporal patterns [30]. Pattern recognition classifiers might allow for cascading power outages to be recognized from Grid Watch data.

## 4.2 Coverage

We recognize that the efficacy of Grid Watch depends on motivating the public to install and run our system. We are optimistic that we would achieve some level of penetration given the high participation rate in several non-monetarily incentivized community-sensing projects, which we discuss in Section 5. In addition to the simple penetration provided by Good Samaritan participants, we aim to add features such as outage statistics, estimated time to power return, and utility comparison that provide sufficient value-add to motivate additional users to join the Grid Watch platform. Several of these such features are provided by existing utility company apps, which have several thousand installed users [6], and further serve our conviction that the potential installed user-base for Grid Watch is large enough to be effective.

For Grid Watch, coverage also refers to Grid Watch’s usage case coverage. While a majority of smartphones are plugged in at night while people are sleeping, the limited set of people who work nights, work at home, or otherwise may leave their phone charging during the day presents an intriguing challenge for Grid Watch. However, the growing sector of ultrabooks and convertible tablets presents another opportunity to expand the coverage of Grid Watch. As per the Windows 8 Ultrabook Specification, convertible tablets must contain an accelerometer to be considered an ultrabook, and regular ultrabooks are recommended to have accelerometers. We believe that microphones are fairly standard across all mobile devices, and regular tablets already contain many of the same sensors that smartphones contain. Even if Grid Watch can only perform widespread characterization of the power grid while people are sleeping at night, we would consider it a success. However, the smaller population segments and other information vectors discussed presents an interesting research direction to push the coverage limits of a crowd-sourced and automated power outage sensing system.

## 4.3 Technical Deployability

Once the initial hurdle of encouraging people to participate in Grid Watch has been surpassed, there are further challenges that come with growing Grid Watch to global scale.

The Android ecosystem largely accommodates Grid Watch, providing us with an easy to access marketplace and a strong API which allows us to differentiate charging sources on Android phones. However, the iOS ecosystem presents us with a few problems. In iOS 6.0, only six types of applications are allowed to run as long-term background programs. This

means that getting a non-developer addition of our Grid Watch app approved by Apple would require either a very flexible reading of the background application requirements, or for Apple to revise their policy regarding background applications. This is a challenge faced by many would-be community sensing applications, such as earthquake monitor or nuclear detectors that are beginning to emerge. We are hopeful, however, given Apple’s new M7 chip and the focus on long-term background data collection using only in-phone sensors, that a new class of Apple-sanctioned applications will emerge.

In addition to software challenges, hardware diversity plagues all application developers. While our limited survey from Figure 2 was able to successfully run common software to extract the 120 Hz peak, we recognize the probability that a greater array of microphones would increase the challenge of ensuring that our AC-mains presence classifier remains effective.

## 4.4 Data Integrity

As Grid Watch begins to accumulate data, it becomes important to develop some metrics to establish the quality of the Grid Watch data. In areas where power companies are well-instrumented and share data, this provides an excellent check. For regions where Grid Watch seeks to supplant utility data, other means of validation are required.

In practice, there are often many other events that can be correlated to a power outage. In Kenya, for example, many customers currently publicly tweet outage reports to the national utility. Other possible avenues include weather reports or newly emerging global Internet health surveys—a geographically clustered area of server outages likely indicates a physical failure of some kind.

Focusing internally, there are other analyses on the Grid Watch data itself that can further provide integrity checks. Existing systems for load forecasting use a diverse array of techniques such as time-series predictors, neural networks, nearest-neighbor approaches, and QP [20] to model the grid. Running these models on our Grid Watch data may provide insight on how well Grid Watch models the grid, how well the models adhere to recorded data, or both.

Finally, we recognize that the Grid Watch system remains entirely vulnerable to “Bad Actors”. It remains an open question as to whether it is necessary to protect Grid Watch from intentional manipulation and if so the correct mechanisms for this protection may be.

## 4.5 Recovery Rate and Data Resolution

The current Grid Watch application is focused on detecting and characterizing power outages. Unfortunately, this misses the perhaps equally interesting characterization of the rate of power outage recovery. A key component of measuring grid health is to evaluate not only how often the grid fails but how quickly and effectively we are capable of repairing it.

One way to do this would be to allow a central service to query sensors on the phone, which combined with context detection and GPS may allow phones to guess if they should

be able to detect AC mains, and then see if they can actually detect it. Furthermore, the ability to perform this type of query would allow for on demand increases in data resolution by using the event detection of one phone to wake up other Grid Watch clients.

## 4.6 Increased Sensor Utilization

The variety of sensors in smartphones raises the possibility of developing new and novel classifiers for detecting outages more reliably. In addition, these sensors could be tasked to further monitor the health (e.g. phase) of an active power grid. For example, the gyroscope and magnetometer could be used to provide greater false positive detection and additional characterization of AC mains frequency.

It has been shown that a 50 Hz<sup>1</sup> fundamental AC frequency can be extracted from digital audio recordings with a high degree of accuracy [22]. This work supports the Grid Watch technique of monitoring for AC mains by analyzing audio captured from the microphone. Currently, power companies use phasor measurement units, or PMUs, to characterize frequency swings and differences in a wide area network [12]. The magnetometers and/or microphones on smartphones may allow Grid Watch to act as a low resolution PMU, allowing us to generate frequency differential maps to further increase visibility into the power grid.

## 5. RELATED WORK

**Grid Health:** Access to the real-time power grid status is critical to its stability. Grid modeling and response is well-studied, but these models require dedicated instruments to gather accurate real-time data of the power state. [15, 17, 23]. Grid Watch aims to provide data to support this analysis with commodity mobile phones.

**Grid Data:** Out of seven United States power companies surveyed<sup>2</sup> none provide historic outage data. These companies do display real-time high resolution outage information on their websites, although this data is not made available in an easily exportable format. The Department of Energy requires utility companies to report outages that affect over 50,000 customers for more than an hour and compiles this data into public annual reports [3].

The World Bank tracks the number of power outages that firms experience in a typical month in countries around the world and makes this data accessible [31]. This information relies on surveys and only reports company level outages. To the best of our knowledge, there exists no automatically updated individual level outage data repository.

**Community Sensing:** Previously deployed community sensing projects have attracted high amounts of participation. As of July 2013, the Zooniverse community sensing platform contained over 800,000 participants across 12 different projects [13]. Other community projects such as Seti@Home [8], GitHub and Government [5], and Folding@Home [4] have also enjoyed great success despite the lack of monetary incentivization. One community smartphone project that has

enjoyed immense success is Waze, a crowdsourced car navigation program with a community of around 50 million users [10]. We believe that his demonstrates that the same sense of community good that drove computer based crowdsourcing projects can also apply to the smartphone community.

**Outage Detection:** A survey of utility companies<sup>3</sup> shows that companies now leverage automated telephone services, online “outage tools”, smartphone applications, and social media sources as means to report outages. However, these methods still rely on customers to report the outage in a timely manner.

Utility companies have the ability to perform measurements over large-scale systems using supervisory control and data acquisition systems and phasor measurement units [2, 7]. In monitoring individual homes, companies still rely in part on traditional meters which require manual recording by employees in the field. In developed countries, an advanced metering infrastructure (AMI) is being deployed [1] that automates power measurements through the use of “Smart Meters”. However, due to cost and privacy concerns, AMI adaptation varies between countries [14, 16]. In the United States overall adoption has reached less than 30% with substantial government support as of mid 2012 [9]. Grid Watch seeks to fill this gap by providing opt-in, automated, fine-grained power information with minimal infrastructure and deployment cost.

Many smart meters use communication back-ends that rely on the power grid [19, 24, 28], making their utility susceptible to grid failures. In contrast to smart meters, Grid Watch is resistant to power grid failures. In the case of a power loss, Grid Watch endpoints have batteries and mobile networks typically have grid-redundant power supplies [11, 25, 26].

## 6. CONCLUSION

We propose Grid Watch, a global, crowd-sourced grid monitoring platform that leverages a simple side-channel available to smartphones—the charger status—coupled with the reliable and independently powered cellular network to provide a simple, free, and easily deployable grid monitoring solution. Our preliminary results show the viability of collecting tightly time-synchronized power state events from heterogeneous phones and operating systems, demonstrating the viability of our key idea. Much remains to be explored, including the challenges of scaling the system, minimizing false positives and ensuring individual privacy and safety while maintaining the authenticity and integrity of the distributed reports. If deployed at scale, Grid Watch could provide unprecedented public data about the global power grid, to the benefit of the public, utility companies and regulators.

<sup>1</sup>This research was conducted in Poland.

<sup>2</sup>DTE Energy, ComED, PG&E, National Power, Duke Energy, XCLE

<sup>3</sup> DTE Energy, ComEd, Duke Energy, National Power, XCEL and PG&E

## 7. REFERENCES

- [1] Advanced Metering Infrastructre. <http://www.ferc.gov/eventcalendar/Files/20070423091846-EPRI%20-%20Advanced%20Metering.pdf>.
- [2] Application of Phasor Measurement Units for Disturbance Recording. [http://www.gedigitalenergy.com/smartgrid/Aug07/PMU\\_disturbance\\_recording.pdf](http://www.gedigitalenergy.com/smartgrid/Aug07/PMU_disturbance_recording.pdf).
- [3] Electric Disturbance Events (OE-417). <http://www.oe.netl.doe.gov/oe417.aspx>.
- [4] Folding@home. <http://folding.stanford.edu/>.
- [5] github and government. <http://government.github.com/>.
- [6] MyDTE energy app. <https://play.google.com/store/apps/details?id=com.dteenergy.mydte>.
- [7] SCADA. [www.powersystem.org/scada](http://www.powersystem.org/scada).
- [8] seti@home. <http://setiathome.berkeley.edu/>.
- [9] Smart meter deployments continue to rise. <http://http://www.eia.gov/todayinenergy/detail.cfm?id=8590#>.
- [10] Waze social gps maps and traffic. <https://play.google.com/store/apps/details?id=com.waze>.
- [11] FCC 07-177 (emergency cell power mandate). <http://njslom.org/FCC-07-177A1.pdf>, 2007.
- [12] M. Adamiak and R. Hunt. Application of phasor measurement units for disturbance recording. [http://www.gedigitalenergy.com/smartgrid/Aug07/PMU\\_disturbance\\_recording.pdf](http://www.gedigitalenergy.com/smartgrid/Aug07/PMU_disturbance_recording.pdf).
- [13] C. S. Alliance, 2013.
- [14] Alliance for Natural Health. Summary of global smart meter legislation. <http://www.anh-europe.org/Time+to+say+NO+to+smart+meters>.
- [15] S. Amin and B. Wollenberg. Toward a smart grid: power delivery for the 21st century. *Power and Energy Magazine, IEEE*, 3(5):34–41, 2005.
- [16] Carlanwickhamhills@gmail.com. Smart metering projects map. <https://maps.google.com/maps/ms?msid=208141621543957618113.0000011362ac6d7d21187&msa=0&ll=11.178402,39.375&spn=172.501588,349.101562&dg=feature>.
- [17] P. Crucitti, V. Latora, and M. Marchiori. A topological analysis of the italian electric power grid. *Physica A: Statistical Mechanics and its Applications*, 338(1-2), 2004. Proceedings of the conference A Nonlinear World: the Real World, 2nd International Conference on Frontier Science.
- [18] European Regulators Group for Electricity & Gas. Status review on regulatory aspects of smart metering (electricity and gas). 2009.
- [19] S. Galli, A. Scaglione, and Z. Wang. Power line communications and the smart grid. In *Smart Grid Communications, 2010 First IEEE International Conference on*, pages 303–308, 2010.
- [20] G. Giannakis, V. Kekator, and N. Gatsis. Monitoring and optimization for smarter power grids: A signal processing perspective. <http://www.eusipco2013.org/tutorials/smartgrid.pdf>.
- [21] S. Guha, A. Moss, J. S. Naor, and B. Schieber. Efficient recovery from power outage.
- [22] M. Kajstura, A. Trawinska, and J. Hebenstreit. Application of the electrical network frequency (ENF) criterion: A case of a digital recording, December 2005.
- [23] R. Kinney, P. Crucitti, R. Albert, and V. Latora. Modeling cascading failures in the north american power grid. *The European Physical Journal B - Condensed Matter and Complex Systems*, 46(1):101–107, 2005.
- [24] S.-W. Luan, J.-H. Teng, S.-Y. Chan, and L.-C. Hwang. Development of a smart power meter for ami based on zigbee communication. In *Power Electronics and Drive Systems, 2009. PEDS 2009. International Conference on*, pages 661–665, 2009.
- [25] Navigant Research. Off-grid power for mobile base stations. <http://www.navigantresearch.com/research/off-grid-power-for-mobile-base-stations>.
- [26] G. O'Reilly, A. Jrad, R. Nagarajan, T. Brown, and S. Conrad. Critical infrastructure analysis of telecom for natural disasters. In *Telecommunications Network Strategy and Planning Symposium, 2006. NETWORKS 2006. 12th International*, pages 1–6, 2006.
- [27] T. Perry. Researcher says better power outage data would reduce frustration and help hurricane recovery. <http://spectrum.ieee.org/tech-talk/energy/the-smarter-grid/researcher-says-better-power-outage-data-would-relieve-fru> Nov 2012.
- [28] SVP Meter Connect. <http://santaclarafreewifi.com/>.
- [29] The Edison Foundation: Institute for Electric Efficiency. Utility-scale smart meter deployments, plans, & proposals. *IEE Report*, 2012.
- [30] Z. Wang, A. Scaglione, and R. Thomas. Power grid vulnerability measures to cascading overload failures. In *Signal Information Processing Association Annual Summit and Conference, 2012 Asia-Pacific*, pages 1–5, 2012.
- [31] E. S. World Bank. Power outages in firms in a typical month (number). <http://data.worldbank.org/indicator/IC.ELC.OUTG/countries/1W?display=default>.