

Grid Watch: Mapping Blackouts with Smart Phones

Noah Klugman[†], Javier Rosa[‡], Pat Pannuto[†],
Matthew Podolsky[‡], William Huang[†], and Prabal Dutta[†]

[†]Computer Science & Engineering Division
University of Michigan
Ann Arbor, MI 48109
{nklugman,ppannuto,wwhuang,prabal}@umich.edu

[‡]Dept of Electrical Engineering & Computer Science
University of California, Berkeley
Berkeley, CA 94720
{javirosa,podolsky}@eecs.berkeley.edu

ABSTRACT

The power grid is one of humanity's most significant engineering undertakings and it is essential in developed and developing nations alike. Currently, transparency into the power grid relies on utility companies and more fine-grained insight is provided by costly smart meter deployments. We claim that greater visibility into power grid conditions can be provided in an inexpensive and crowd-sourced manner independent of utility companies 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 that reduce the likelihood of false power outage reports, and corroborating actual reports with other phones through data aggregation in the cloud. The proposed 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 utility companies accountable for poor power quality. This paper demonstrates the viability of the basic idea, identifies a number of challenges that are specific to this application as well as ones that are common to many crowd-sourced applications, and highlights some improvements to smartphone operating systems that could better support such applications in the future.

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, regulators, and ratepayers. 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 approach can offer a detailed view into power grid conditions, the approach also has its drawbacks. As smart meters are under the aegis of 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 crowd-sourced method of characterizing power grid conditions. The key insight underlying Grid Watch is that smartphones can easily detect power outages by monitoring changes in their charge state when plugged into the wall. Furthermore, smartphones have various sensors which allow them to locally verify that a change in power state while charging is a likely power outage rather than a manual disconnection from the charger. For example, a microphone can detect the presence or absence of an ambient 50 or 60 Hz "hum," or its harmonics, which provides some evidence of the presence or absence of AC electric fields emanating from nearby power lines.

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 slow. Finally, Grid Watch can augment utilities' existing monitoring systems, potentially providing visibility into locations that the utilities are unable to monitor easily, helping them isolate and correct problems in their own networks.

To the best of our knowledge, there is no publicly available repository of power outage data, much less one that is automatically updated independently of utility-generated reports. We believe Grid Watch could play an integral role in the creation of such a data set, enabling third parties – ratepayers, researchers, and regulators – to analyze power grids and hold utilities accountable.

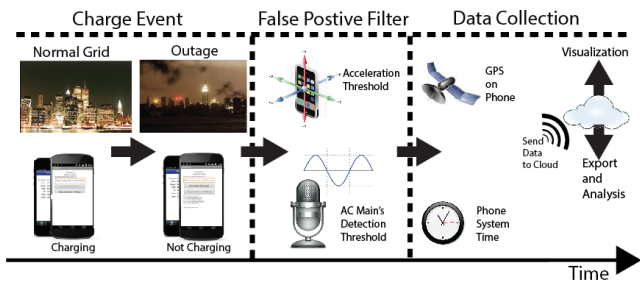


Figure 1: Grid Watch operation diagram. A plugged-in phone changes from a powered state to an unpowered state with grid failure. Grid Watch registers this event, verifies that it is not a likely false positive, and reports the event to the cloud for analysis, export and visualization.

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 rarely 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 for export and visualization and is used as input data into grid behavior modeling algorithms. An overview of the system is shown in Figure 1. We implement 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 wake up the Grid Watch app from the background. When an OS event registers that a phone has stopped charging, Grid Watch briefly samples the accelerometer and microphone (currently 5 seconds). The accelerometer detects if the phone is being moved (unplugged), and an FFT on the audio samples detects the AC mains hum. In addition, Android’s API exposes the classification of charger type, allowing Grid Watch to 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 are not automatically detected by Grid Watch (e.g. an outage that occurs when the user’s phone is not plugged in) and reclassify falsely reported outages.

2.2 The Data

The current Grid Watch implementation collects the following data, which we consider to be the minimum needed for Grid Watch to be effective:

GPS Location. 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 (e.g. truncated at the house, block, neighborhood, or city granularity). 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 and preserve user privacy.

Classifier Results. Both of our local outage filters are threshold-based: 1) Did the accelerometer move “too much”? 2) 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 ground truth of when an outage is observed. In addition, the Grid Watch service records a timestamp when the event is actually received. A high delta between this timestamp and the timestamp from the central service when the event is received could suggest the loss of independently powered cellular infrastructure.

User Data. Additional information from the user is gathered optionally by Grid Watch independent of an outage event. This information includes the name and contact information for the utility company that provides power for the user, an estimation of the frequency of power outages that the user experiences, and basic information about the user’s phone charging habits.

Unique ID. An ID is not strictly necessary for the operation of a Grid Watch. We collect it, however, for the purpose of estimating the Grid Watch user base and density in a given area. In addition, we believe users may have an interest in tracking their own 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.

2.3 The Central Service

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

3. EVALUATION

Although we have not yet deployed Grid Watch, we perform a series of experiments to validate several key hypotheses. 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 perform two experiments to validate our methods of false positive detection. First, we run 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, the phones are used routinely. The central service did not receive any false positives during this experiment.

Additionally, we test the ability of phones to detect the hum from AC mains. We record a five second audio sample at a 44.1 Ksps sample rate from four different phones inside a house and turn off the master circuit breaker to simulate a power outage. These recordings contain minimal background noise. We perform an FFT on these recordings to search for the AC mains frequency. The results of this survey are shown in Figure 2. We observe the 60 Hz (U.S. AC mains) frequency clearly from audio recorded on the phones running Android when the power in the house is on. Within the audio recorded on the iPhones, we observe no 60 Hz peak, suggesting 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 the audio recorded on the Android phones. We build our AC mains classifier to detect the presence of a 120 Hz peak. To ensure accurate detection in other countries, our classifier can be set to report the presence of mains power if a 100 Hz peak, the second harmonic of a 50 Hz mains, is detected.

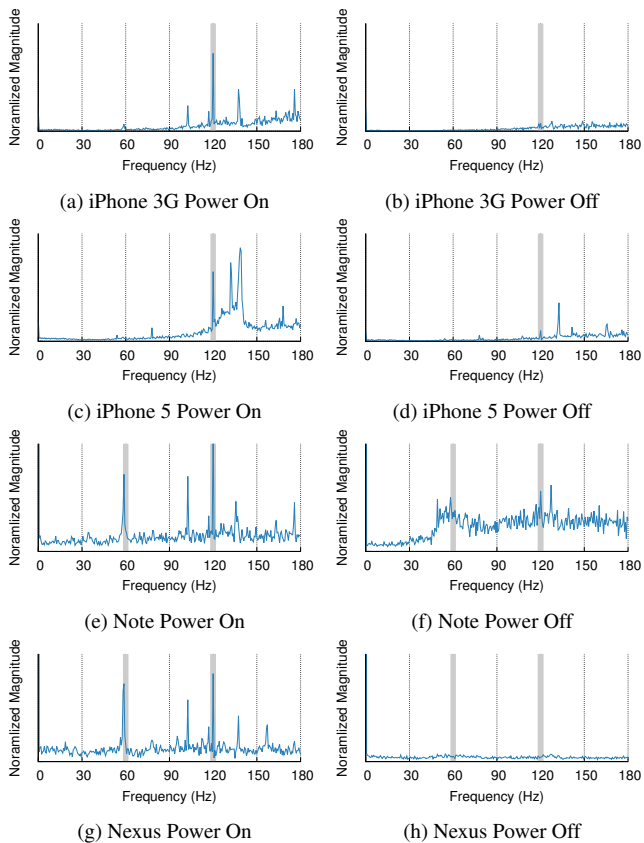


Figure 2: FFT of audio samples captured on four different phones illustrating the absence and presence, respectively, of 60 Hz and its harmonics from ambient electrical fields. The graphs on the left capture the home environment during normal power a power outage for four different phones. On all models, 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 in their audio frontends to presumably mitigate noise from the environment, but none of the harmonics are filtered, allowing our microphone-based detection to remain effective.

We test the ability of a Galaxy Nexus, Nexus 7, and iPhone 5 to detect 60 Hz or its harmonics in different scenarios. We use four different physical environments: 1) A laboratory with a higher than normal concentration of electronic equipment at the University of Michigan, 2) A one bedroom apartment, 3) A house, and 4) An office with a normal concentration of electric equipment at the University of Michigan. In each environment, we test four different types of background noise: 1) No noise, 2) Regular talking, 3) TV on at regular volume, and 4) Music playing at normal volume. Recordings are taken from the phones and AC mains presence is determined visually by inspecting the FFT. The results of these experiments are shown in Figure 2(i).

We find that AC mains is detectable across many environments and background noises. We notice at times both the Nexus 7 and iPhone 5 do not detect any harmonic of AC mains when power is present. Both phones record audio at a significantly higher gain than the Galaxy Nexus suggesting that they might be employing automatic gain control that could be masking the AC mains signal.

Cascading outages spread within minutes [19], setting an upper delay bound for Grid Watch reports to be useful in tracking the spatial and temporal spread of certain outages. To test the temporal clustering of observations of the same power outage event, we connect ten smartphones of six different models to a single power strip. We turn off the power switch, and examine the timestamps each phone generates upon detecting the power outage. We repeat this process 20 times, and display the results in Figure 3.

		Mode	Device	Apt	Home	Lab	Office
POWER ON	Quiet	Galaxy Nexus	TP	TP	TP	TP	TP
		Nexus 7	TP	TP	TP	TP	TP
		iPhone 5	TP	TP	TP	TP	TP
	Talking	Galaxy Nexus	TP	TP	TP	TP	TP
		Nexus 7	TP	TP	TP	TP	FN
		iPhone 5	TP	TP	TP	TP	TP
	TV	Galaxy Nexus	TP	TP	TP	TP	TP
		Nexus 7	TP	FN	TP	TP	TP
		iPhone 5	TP	TP	TP	TP	TP
	Music	Galaxy Nexus	FN	TP	TP	TP	TP
		Nexus 7	FN	TP	TP	TP	TP
		iPhone 5	TP	FN	TP	TP	TP
POWER OFF	Quiet	Galaxy Nexus	TN	TN			
		Nexus 7	TN	TN			
		iPhone 5	TN	TN			
	Talking	Galaxy Nexus	TN	FP			
		Nexus 7	FP	TN			
		iPhone 5	FP	TN			

(i) Results of our AC mains peak detection classifier on a variety of phones in a variety of environments. Power off in the Apt environment cut power in the apartment, but not the rest of the building. We are unable to cut power to the lab or office. Our AC mains classifier shows 87% accuracy in our limited sample, which is encouraging but clearly indicates that smartphone-based detection alone is not sufficient for low error rates. True positive (TP) and true negative (TN) indicate either the detection or lack of detection of AC mains when expected. Conversely, false positive (FP) and false negative (FN) indicate either the detection or lack of detection of AC mains when not expected.

We find that the average standard deviation between reported times for an identical power outage was 0.76 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 0.8 s. The maximum time difference between timestamps for a single event is 6.21 s. The maximum time difference is less than 2.2 s.


Based on this data, 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. Finally, for WiFi connected smartphones, we may be able to employ NTP.

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 for our particular application of grid monitoring. In this section we enumerate what we see as the greatest challenges moving forward with Grid Watch and similar community sensing efforts.

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 [17, 28]. As a result, customer provided reports supply the most actionable data but



(a) Ten phones detecting the same twenty simulated power outages

Figure 3: Exploring how tightly-coupled timing event detections are between a variety of phones. For each trial, phones are connected to the same power strip, which we switch off to simulate a power outage. The phones time-stamp the event detection, and we characterize the standard deviation in reported times for each trial. Only the charging state classifier is active during these experiments. We find that across a diverse array of phones and chargers, the reported event time is within one second in 19 of 20 trials.

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 [19, 20]. 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 are 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 [29]. 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 support our claim that the potential installed user-base for Grid Watch is large enough to be effective.

Coverage also refers to Grid Watch’s use 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 that contain accelerometers present another opportunity to expand Grid Watch coverage. Microphones are common 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 present an interesting research direction to push the coverage limits of a crowd-sourced and automated power outage sensing system.

Additionally, we envision older phones that are no longer in use being repurposed as stationary Grid Watch instances. Assuming these phones can still access the cellular infrastructure, they would be free from many of the constraints in coverage that arise with an everyday-use phone.

4.3 Platform Challenges

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, none of which describe the Grid Watch app. This means that getting the Grid Watch app approved by Apple for deployment 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 facing many would-be community sensing applications, such as earthquake monitoring or nuclear detectors that are now emerging. 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 shows that the 120 Hz peak can be extracted from audio recorded on a subset of phones, 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 will become important to develop 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 necessary and must be devised.

In practice, there are often many other events that can be correlated to a power outage. In Kenya, for example, many customers 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 methods of analysis that can be performed on the Grid Watch data itself to further check integrity. Existing load forecasting systems use a diverse array of techniques such as time-series predictors, neural networks, nearest-neighbor approaches, and QP [19] 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 vulnerable to “bad actors” who “bear false witness”. It remains an open question as to whether it is necessary to protect Grid Watch from intentional manipulation and if so what the correct mechanisms for this protection may be.

4.5 Recovery Rate

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 it is repaired.

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 the phones 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. This mechanism could also help corroborate true witnesses.

4.6 Data Resolution and User Privacy

There is possibly some concern to user privacy when Grid Watch reports the location of an outage. Publicly available Grid Watch data could betray homes that have lost power and in turn their burglary detection systems. High-fidelity data could potentially reveal individual Grid Watch users.

Currently, our Grid Watch implementation records both a high-precision GPS location and a “low-precision” network-based location (e.g. cell tower or nearest WiFi AP, depending on platform APIs). Once we have collected enough Grid Watch data, we intend to explore the required fidelity for Grid Watch to be effective such that we minimize the invasiveness of Grid Watch on user privacy. In the meantime, the Grid Watch application provides a mechanism for users to limit the fidelity of the reported location data by truncating GPS coordinates, effectively limiting its precision.

4.7 Increased Sensor Utilization

The variety of sensors in smartphones raises the possibility of developing new and novel classifiers for detecting outages more reliably. Furthermore, heuristics such as the presence of WiFi signals or classification of captured sound (e.g. is there music playing?) around the phone are two examples of additional techniques that might decrease false positive reporting.

In addition, these sensors could be tasked to monitor the health (e.g. phase) of an active power grid. Currently, power companies use phasor measurement units, or PMUs, to characterize grid phase and frequency in the 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 or detect sudden phase changes to further increase visibility into the power grid.

An additional measure of grid stability is the frequency deviation from AC mains at different locations. It has been shown that a 50 Hz¹ fundamental AC frequency can be extracted from digital audio recordings with a high degree of accuracy [22]. While that work was sufficient for post-hoc fingerprinting, it remains to be seen whether the audio channel can provide real-time monitoring. Because of the trade-off between temporal and frequency resolution inherent to an FFT², it is unclear if we can achieve a high enough resolution frequency measurement from the phone to be useful, limited both by the data rate available from the audio front-end and the energy cost of performing high-resolution FFTs.

¹This research was conducted in Poland.

² For example, when sampling at 44.1 kHz for 5 seconds we gather 220,500 samples giving us $\frac{220,500}{2} = 110,250$ bins. Our bin resolution is then limited to $\frac{44.1 \text{ kHz}}{110250} = 0.4 \frac{\text{Hz}}{\text{bin}}$.

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. [14, 16, 23]. Grid Watch aims to provide data to support this analysis with commodity mobile phones.

Grid Data. Out of seven United States power companies surveyed³ none provide long term 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 [30]. 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 [15]. Other community projects such as Folding@Home [4], GitHub and Government [5], and SETI@Home [8] have also enjoyed great success despite the lack of monetary incentives. 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].

Outage Detection. A survey of the same utility companies 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 adoption varies between countries [13, 21]. 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 [18, 24, 27], 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 backup power supplies like battery banks and generators. [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

³ DTE Energy, ComEd, PG&E, National Power, Duke Energy, and XCEL.

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 researchers, ratepayers, and regulators.

7. ACKNOWLEDGMENTS

This material is based upon work supported by the Development Impact Lab (USAID Cooperative Agreements AID-OAA-A-13-00002 and AID-OAA-A-12-00011), US Agency for International Development. This work is supported by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA. This material is based upon work supported by National Science Foundation grant #1059372 (“CI-ADDO-NEW”). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

The authors would also like to thank the reviewers for their insightful comments and feedback and the editing and support provided by Brad Campbell and Zakir Durumeric.

8. REFERENCES

- [1] Advanced metering infrastructure. <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.
- [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.
- [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.
- [13] Alliance for natural health. Summary of global smart meter legislation. <http://www.anh-europe.org/Time+to+say+NO+to+smart+meters>.
- [14] 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.
- [15] Citizen Science Alliance. Zooniverse. <https://www.zooniverse.org/>.
- [16] 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.
- [17] European Regulators Group for Electricity & Gas. Status review on regulatory aspects of smart metering (electricity and gas). 2009.
- [18] 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.
- [19] 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>.
- [20] S. Guha, A. Moss, J. S. Naor, and B. Schieber. Efficient recovery from power outage.
- [21] S. Harrison. 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>.
- [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] SVP Meter Connect. <http://santaclarafreewifi.com/>.
- [28] The Edison Foundation: Institute for Electric Efficiency. Utility-scale smart meter deployments, plans, & proposals. *IEE Report*, 2012.
- [29] 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.
- [30] E. S. World Bank. Power outages in firms in a typical month. <http://data.worldbank.org/indicator/IC.ELC.OUTG/countries/1W?display=default>.