Hardware, Apps, and Surveys at Scale: Insights from Measuring Grid Reliability in Accra, Ghana

Noah Klugman
University of California, Berkeley
nklugman@berkeley.edu

Joshua Adkins
University of California, Berkeley
adkins@berkeley.edu

Susanna Berkouwer
University of California, Berkeley
berkouwer@berkeley.edu

Kwame Abrokwa
Brixels Company Limited, Accra
kwame.abrokwa@gmail.com

Ivan Bobashev
University of California, Berkeley
ivan_bobashev@berkeley.edu

Pat Pannuto
University of California, Berkeley
ppannuto@berkeley.edu

Matthew Podolsky
University of California, Berkeley
podolsky@berkeley.edu

Aldo Suseno
University of California, Berkeley
aldosuseno@berkeley.edu

Revati Thatte
University of California, Berkeley
revati.thatte@berkeley.edu

Catherine Wolfram
University of California, Berkeley
cwolfram@berkeley.edu

Jay Taneja
University of Massachusetts, Amherst
jtaneja@umass.edu

Prabal Dutta
University of California, Berkeley
prabal@berkeley.edu

ABSTRACT

The vision of sensor systems that collect critical and previously ungathered information about the world is often only realized when sensors, students, and subjects move outside the academic laboratory. However, deployments at even the smallest scales introduce complexities and risks that can be difficult for a research team to anticipate. Over the past year, our interdisciplinary team of engineers and economists has been designing, deploying, and operating a large sensor network in Accra, Ghana that measures power outages and quality at households and firms. This network consists of 457 custom sensors, over 3,000 mobile app instances, thousands of participant surveys, and custom user incentive and deployment management systems. In part, this deployment supports an evaluation of the impacts of investments in the grid on reliability and the subsequent effects of improvements in reliability on socioeconomic well-being. We report our experiences as we move from performing small pilot deployments to our current scale, attempting to identify the pain points at each stage of the deployment. Finally, we extract high-level observations and lessons learned from our deployment activities, which we wish we had originally known when forecasting budgets, human resources, and project timelines. These insights will be critical as we look toward scaling our deployment to the entire city of Accra and beyond, and we hope that they will encourage and support other researchers looking to measure highly granular information about our world's critical systems.

CCS CONCEPTS

• **Hardware** → **Sensor applications and deployments; Energy metering**; • **Social and professional topics** → **Systems development; System management**.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

COMPASS '19, July 3–5, 2019, Accra, Ghana

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6714-1/19/07.

<https://doi.org/10.1145/3314344.3332482>

ACM Reference Format:

Noah Klugman, Joshua Adkins, Susanna Berkouwer, Kwame Abrokwa, Ivan Bobashev, Patrick Pannuto, Matthew Podolsky, Aldo Suseno, Revati Thatte, Catherine Wolfram, Jay Taneja, and Prabal Dutta. 2019. Hardware, Apps, and Surveys at Scale: Insights from Measuring Grid Reliability in Accra, Ghana. In *ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS) (COMPASS '19)*, July 3–5, 2019, Accra, Ghana. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3314344.3332482>

1 INTRODUCTION

The power grid is arguably the most complicated machine humanity has built, and the payoffs from this marvel have been transformative. No country has achieved economic industrialization without significant increases in energy use. Hospitals, schools, factories, and homes across the world depend on electricity for their daily operations. As such, the developing world has seen tremendous investments in the electricity grid in recent years.

Investments in the electricity sector in the developing world tend to focus on increasing *access* to electricity by expanding the reach of the grid. There has been less focus on the *quality* of that access. Improvements in electricity reliability can be harder to measure—and to achieve—than improvements in access. Many factors influence power quality. To improve reliability, a utility needs fine-grained information about how different attributes of the grid perform in addition to a well-developed toolkit to address problems [4]. Many electrical utilities are already under-resourced for the enormously complex and expensive task of planning, extending, and operating their current systems, and they cannot easily launch programs to collect the data needed for increasing reliability [27]. Neglecting reliability, however, is often associated with a reduction in the demand, utilization, and social benefit of electricity [13].

Recognizing this, investments are increasingly aimed at improving the *reliability* of electricity distributed on the existing grid. In 2014, the Millennium Challenge Corporation (MCC) and the Government of Ghana signed the Ghana Power Compact, a USD 498 million investment designed to improve the grid generation, transmission, and distribution systems in Ghana to be implemented by the newly created Millennium Development Authority (MiDA) [3].

As independent evaluators of the Compact, we have been working with these partners to design and deploy sensors that can measure power outages, voltage fluctuations, and frequency instabilities at the low-voltage level of the distribution grid in Accra, Ghana. Our interdisciplinary team consists of economists evaluating the socio-economic impacts of these investments and engineers building a system that can provide the data requirements of that evaluation.

This paper reports on both the sensing methodology developed and our experiences during the first year of pilot deployment activities in several districts of Accra, Ghana. We deployed our system at three different scales, with the final iteration resulting in a sensor system that is currently collecting the highest resolution information on low-voltage power reliability in the city of Accra. We deployed 362 custom sensors (called PowerWatch) that are deployed at households and firms, recruited over 3,000 participants to download a mobile app (called DumsorWatch) that attempts to sense power reliability issues, and surveyed nearly 3,000 participants.

We discuss the technology deployed, the design of our deployment, and where our planning and assumptions failed or caused unexpected problems. We attempt to categorize our key challenges and describe the steps that we have taken or will take to overcome each of these challenges. We find that each level of scale brings unique complexities for both engineering and operational tasks. While some of these complexities are one-time costs, many can be attributed to the continuous nature of operating and managing a sensor deployment at scale. This combination of scale and continuity stretches the administrative ability of the university system, explodes the amount of state that must be kept to manage the sensors, amplifies errors in data collection, and ultimately requires the development of automated tools to facilitate tasks that our field team and ourselves could not handle at scale. It is our hope that these lessons will inform future efforts to deploy continuous monitoring and evaluation systems in developing regions.

2 DEPLOYMENT GOALS

This paper focuses on our experiences designing, implementing, and iterating a deployment methodology at three different scales in Accra, Ghana. In order to provide context for how this methodology evolved, we begin by introducing the research questions the deployment was designed to address, the results of which are out of scope for this work.

2.1 Exploring Impacts of Reliability

The causal relationship between electricity reliability and socioeconomic well-being is not well understood. Anecdotally, frequent outages constrain economic well-being by reducing the benefits from welfare-improving appliances like fans and refrigerators or income-generating assets like sewing machines. The deployment was in part designed to generate reliability and socioeconomic data for an ongoing economic study that exploits two quasi-random sources of variation in reliability in Accra. By comparing households and firms whose socioeconomic characteristics are identical in expectation, and that differ *only* in terms of the quality and reliability of power they receive, we will be able to estimate the causal effect of these attributes on socioeconomic outcomes such as well-being, productivity, and health for the residents of Accra.

Figure 1: Overview of deployment. To support the goals of the deployment, our team selects sites that are being improved by the Ghana Power Compact and control sites. The technology is deployed in both sites along with surveys at the beginning and end of the deployment. This lets us meet our goals of evaluating the impact of grid improvements to power reliability and the socioeconomic impact of that reliability on consumers.

2.2 Improving Energy Reliability Data Quality

Two common metrics of energy reliability are the System Average Interruption Duration Index (SAIDI) and the System Average Interruption Frequency Index (SAIFI) [1]. These are also key performance indicators for the Ghana Power Compact. The construction of SAIDI and SAIFI is shown as Equation (1) and Equation (2).

$$SAIDI = \frac{\text{Total duration of sustained interruptions in a year}}{\text{Total number of consumers}} \quad (1)$$

$$SAIFI = \frac{\text{Total number of sustained interruptions in a year}}{\text{Total number of consumers}} \quad (2)$$

Currently, the Electric Company of Ghana (ECG), the power utility operating in Accra, depends on customer calls to estimate the numerator for both SAIDI and SAIFI at the low-voltage level, and uses a supervisory control and data acquisition (SCADA) system that contains sensors on feeder lines, substations, and transmission lines to estimate the numerator for outages that occur at medium and high voltages. The resolution of data from customer calls suffers from a number of problems, including that customers may not call, customers may not experience all outages (i.e., while they are sleeping), and few customers would call on restoration (making duration for SAIDI difficult to estimate).

Our deployment aims to improve the estimation of the numerator of both SAIDI and SAIFI by placing sensors in the field that automatically report the location and duration of power outages.

2.3 Developing an Independent Measurement Methodology

It is understood how to measure power outages with large-scale deployments of expensive supporting infrastructure like smart meters and traditional SCADA systems. It is less understood how to take these same measurements when one cannot depend on utility participation, cannot—due to either monetary or time restrictions—implement large-scale infrastructure deployments, and may require the agility to conduct deployments in multiple regions. This deployment was designed to evaluate the feasibility of a novel and agile methodology for deploying sensors independent of the utility and the efficacy of this methodology for measuring energy reliability.

3 DATA COLLECTION INSTRUMENTS

We developed two types of data collection instruments, sensors and surveys, to achieve the goals described in Section 2. These instruments both collect the data to inform our results and enable a measurement methodology that is independent of the utility.

3.1 Sensors

We developed two different sensors that detect the presence and absence of grid power: an app called DumsorWatch that is installed on a participant's mobile phone, and a fixed-point sensor called PowerWatch that is plugged in at a household or firm.

3.1.1 DumsorWatch. Our mobile sensing technology in Ghana is called DumsorWatch, and is an Android app that is installed on the everyday-use smartphone of a participant who lives and/or works in Accra. The “Dumsor” in the app name is the local word for power outages meaning “off on” and was chosen for stronger branding and association with power outages in the Ghanaian context. DumsorWatch automatically senses power outages and power restorations by using a combination of on-phone sensors and cloud services and is based on prior work [17]. If a phone can't access the Internet when an outage is sensed by DumsorWatch, the app will queue data to be sent when connectivity is restored.

3.1.2 PowerWatch. The fixed-point sensing technology we designed and deployed is called PowerWatch. PowerWatch plugs into a participant's home or firm, and integrates power reliability sensors with a GSM radio, allowing for measurements to be sent in near-real time to a cloud service. Plugging into a participant's home or firm as opposed to directly measuring the electric grid means we do not need prior approval or cooperation from the utility to deploy the sensors, a primary goal of our deployment described in Section 2.3. PowerWatch senses power outages and power restorations timestamped to the millisecond, GPS-based location, voltage, and grid frequency. PowerWatch contains a battery to allow for continuous reporting throughout a power outage and will queue data if there are connectivity problems with the GSM network.

3.2 Surveys

A socioeconomic survey of approximately 60 minutes in length accompanied the deployment of each PowerWatch device with a respondent, and a shorter survey was administered to respondents who did not receive PowerWatch but did download DumsorWatch. All surveys were completed using SurveyCTO and participants were incentivized for their time. Surveys were verified using high-frequency checks to address any obvious data quality issues. Example data collected includes:

- (1) Demographics: name, age, education, income.
- (2) Electricity attributes: appliance and surge protector ownership, usage of electricity and generators.
- (3) Recall of power quality in the past 2, 7, and 30 days.
- (4) Social media usage and perceptions of the energy crisis.

Along with providing data, the survey was used to support the deployment and development of the technology itself. For example, we recorded in the survey a unique code for the PowerWatch device and DumsorWatch app deployed with each respondent, and their phone number and GPS location, so that the sensors could later be

Figure 2: Deployment methodology of sensors. By randomly sampling households and firms under a transformer, sensors are capable of sensing high voltage (HV), medium voltage (MV), and a significant portion of low voltage (LV) outages. Sensors may not detect single phase outages because our sampling does not guarantee sensors are distributed across all possible phases in practice, as in the bottom outage of (d). This is due to both the difficulty of identifying the phase(s) to which a service is connected and manual phase switching by a household or firm. Sensors estimate the average frequency and duration of outages, which include both single-phase and service-level outages.

associated to individuals. To inform DumsorWatch debugging, we asked about the way that residents of Accra employ their mobile phones, how many phones and SIM-cards they use, and how frequently they upgrade their phones. To inform the deployment of the PowerWatch device, we recorded whether the respondent turns off their electricity mains at night and whether they had any safety concerns about PowerWatch.

4 DEPLOYMENT METHODOLOGY

We design a deployment methodology to achieve the goals described in Section 2. Our methodology deploys the data collection instruments at specific locations on the grid in order to both monitor the success of grid improvements performed in the Ghana Power Compact and compare socioeconomic indicators across differing levels of reliability. We also design and deploy deployment management tools to assist in our implementation of the methodology.

While we describe multiple deployments in this paper, each with differing levels of scale and evolving goals, the overall structure of the deployment methodology remained consistent across these deployments. First, we develop criteria for site selection that allow us to answer specific socioeconomic questions. Then, we devise a sampling scheme that gives sufficient coverage of each chosen site as well as sufficient redundancy to enable cross-validation of the new measurement technology. We then work with a team of field officers to deploy in the chosen regions, employing deployment management tools to maintain and monitor the system. The rest of this section considers each of these components in detail.

4.1 Site Selection

We select a subset of the sites where infrastructure upgrades are planned (*'treatment sites'*) and then quasi-randomly select a set of sites that are comparable in observable characteristics (*'control sites'*). For each site, we then define a geographic surveying area that is the intersection of a 200 meter radius from the centroid of the site, and a 25 meter region extending from the low-voltage network being measured. This analysis is performed using GIS tools operating on a map of the grid in Accra provided to us by an independent contractor implementing the grid improvement work.

Once the specific sites are selected, we target a deployment of three PowerWatch devices and 20 DumsorWatch app downloads at each site. Using the GIS technology described above, we produce a series of maps marking the geographic area bounding each site. Field officers use these maps, along with the GPS coordinates for the sites, to identify the surveying area and deploy sensors accordingly.

4.2 Sampling Strategy

We deploy our sensors with residents of Accra, either at their home or place of work (or both, if these are co-located), with an attempted 50% split between households and firms. Installing PowerWatch at consumer plugs and DumsorWatch on consumer phones allows us to not depend on direct access to utility infrastructure such as transformers or lines, and to measure power quality without utility participation at the point where it is least understood: the customer.

Our strategy is built around redundant sampling such that multiple sensors are placed under a single transformer. When all sensors in this group report an outage at the same time, we can be confident it was due to an issue affecting the transformer rather than a single customer. Further, when we observe sensors below multiple different transformers reporting outages simultaneously, we can infer the outage occurred at a higher level of the grid. This sampling strategy is shown in Figure 2.

4.3 Deployment and Surveying Team

We develop a local staff structure that, compared to traditional survey work, uniquely supports our continuously operating deployment. This involves employing a full-time local field manager to oversee the initial roll-out and on-going maintenance of the system, and an auditor to follow up with participants who report problems over the phone and with sensors no longer functioning.

To implement our medium and large scale deployments, we temporarily employ an additional team of 10 field officers and three team leads. Field officers find potential participants, screen their eligibility, and get informed consent. They then perform the survey, install the sensors, and answer any participant questions. We conduct multiple training exercises with the entire team where each member learns about the technologies being deployed, and practices completing the survey and deploying the technologies.

Field officers visit sites in groups of two to alleviate safety concerns. We provide team uniforms to make it clear they are part of an official project, as shown in Figure 4. We also provide backpacks to carry supplies, tablets for the survey, WiFi hotspots to upload the survey and download the DumsorWatch app, flashlights for safety, and feature phones to verify the phone numbers of participants to ensure we know where to send the participation incentives.

4.4 Depending on Participants

The placement of PowerWatch sensors directly in homes and firms—where participants can unplug them, run generators, or fail to pay their meter—increases the noise of our data relative to a deployment on utility-owned equipment such as transformers. Similarly, the DumsorWatch app can be uninstalled, reducing coverage and leading to a potentially under-sampled signal.

In a preemptive attempt to decrease the noise caused by a sampling strategy screen participants for specific criteria including owning a phone with Android version 4.1–8.1 and being an active customer on the grid. We explain the goals, risks, and benefits of the project, and seek written consent. Finally, we provide a phone number if participants have any questions or concerns.

To further encourage continuous participation, we compensate participants monthly with airtime credits on their mobile phone. Participants receive a small amount of airtime for initial recruitment, 4 Ghana Cedi (0.75 USD) monthly for keeping DumsorWatch installed, and 5 Ghana Cedi (0.93 USD) monthly for keeping PowerWatch installed. These numbers were settled on in consultation with our field manager. Additionally, participants who have a PowerWatch sensor placed at an outlet in their home receive a power strip so that the sensor does not take up a needed outlet.

4.5 Deployment Management Tools

We developed three software subsystems to support the deployment. These include an automated incentive system to transfer the airtime incentives, a deployment management system to a) keep track of sensor and participant state and b) display deployment health to the field management team, and a data visualization and analysis system. These were developed as a result of our experiences as the development scaled over time and are discussed, including the specific experiences that lead to their inception, in Section 6.1.

5 EXPERIENCE

In this section we share our experiences deploying in Accra, emphasizing those that were more complex and/or costly than we originally forecast. Our experiences differed depending on the scale of the deployment we were running; each scale uncovered its own complexities. These scale-specific challenges are presented as such, along with their root cause and our mitigation strategy. These complexities are captured in Table 1.

5.1 Overview

Between May 2018 and June 2019 we completed three deployments in Accra at three different scales: a small-scale pilot, a medium-scale deployment, and a large-scale deployment. To date, 3,400 individuals in Ghana have downloaded the DumsorWatch app and we have installed 457 PowerWatch sensors, monitoring 151 transformers across ECG's Achimota, Kaneshie, and Dansoman districts in Accra. For the medium- and large-scale deployments, we incentivized all participants and completed a baseline recruitment survey. A subset of participants, including all participants with a PowerWatch device, completed a longer socioeconomic survey. In December 2018, we conducted a follow-up survey with 462 respondents to understand their participant experiences and to collect updated measures of time-varying socioeconomic outcomes.

Figure 3: Evolution of PowerWatch with each deployment. PowerWatch revision A consisted of an off-the-shelf compute/communication module and enclosure (A.1) and paired with a custom sensor front-end (A.2). Data from this revision informed the need for a better enclosure and more casing in revision B, which consisted of a custom sensing and communication board (B.1), enclosure with externally plugged power supply (B.2), and a separate grid voltage and frequency sensor (B.3). While the separate grid voltage and frequency sensor allowed for easier assembly, its complications led us to build revision C, a completely encased custom sensor which plugs directly into the wall, to sense grid voltage and frequency.

5.2 Small-Scale Pilot

The first activity we performed was a deployment of 15 PowerWatch sensors and 5 DumsorWatch apps. The goal of this deployment was to validate that the technology can reliably sense power outages and transmit this information over many weeks in the field. We performed no survey work and no site selection work for the small-scale pilot. The primary challenges were related to producing the technology, connecting the PowerWatch sensors to the cellular network, and building enough local capacity for PowerWatch and DumsorWatch to be deployed.

In addition to testing the technology, we worked to build relationships to support future scaling. We reached out to the local stakeholders to get feedback on the assumptions driving the designs of the sensors used in our deployment, and were able to speak with engineers and managers at ECG, the Millennium Development Authority (MiDA), and various independent contractors involved in the Ghana Power Compact. Further, we received data from ECG that helped validate our hypothesis that the measurements of SAIDI and SAIIFI could benefit from higher resolution measurements.

Even at small scale, we experienced unanticipated technical challenges. To get PowerWatch on the cellular network we initially attempted to use the global Particle IoT SIM cards which were included with our cellular modems. We found its connection to be much spottier than a local SIM, and Particle support had little introspection into the problem. Because of this, we decided to use SIM cards from the largest local carrier (MTN), but we encountered a 3 SIM-card-per-person limit upon purchase. Although we were able to get around this by visiting different stores, purchasing SIM cards in stores was not an option for future scale.

Another challenge was keeping SIM cards functional. Prepaid SIM cards require the purchase of data plans for the SIM, which is done using a USSD application that can only be run from within Ghana; there is no web-based account management or top-up available. We initially solved this problem by purchasing a 90-day data plan, the longest available. This was sufficient for our small-scale pilot but would not be viable for future deployments.

5.3 Medium-Scale Deployment

In our medium-scale deployment, 1,981 individuals downloaded the DumsorWatch app and 165 individuals installed PowerWatch sensors. Deployment activities took around one week for training, two weeks to survey participants and deploy PowerWatch sensors, and then another three weeks to conduct short surveys and install DumsorWatch apps. We ran this deployment for seven months.

Unlike the small-scale deployment, this scale required implementing our full deployment design, including hiring a full local implementing team, recruiting and incentivizing participants, choosing deployment sites, extracting value from the data streams, and fully implementing the survey instruments. We enumerate the changes experienced as we increased from small- to medium- scale, paying particular attention to the challenges extracted in [Table 1](#).

5.3.1 Organizational. The medium-scale deployment was large enough that the financial responsibilities were significant. We had to start managing multiple monthly payments for cloud services and payments to local companies for cell network connectivity and incentive transfers. Most of this increase in complexity was ultimately handled by staff at the University of California, Berkeley, but establishing payment schedules took a large amount of effort from the research team. The university still missed payments, causing frequent delays, especially when payment was needed in a short time frame (1-2 weeks).

Because prepaid SIM cards were not viable or purchasable at the quantities now needed, we had to enter into a contract with the cellular provider, MTN. Despite multiple meetings, MTN was initially quite hesitant to provide the SIM cards due to concerns about whether our application was legitimate. We were ultimately able to overcome these concerns by visiting the MTN main office in our university shirts, giving a technical demo, and answering questions about our backgrounds and affiliations.

At this scale, many of the cloud-based software services our systems were built upon were no longer eligible for free-tier usage. For one service in particular, this meant that we were going to be unable to continue with this technology without signing a multi-year contract that extended beyond the length of the deployment. We found a workaround for this deployment by applying to a special program within the company, but in future deployments we would more carefully consider pricing models for ancillary services.

5.3.2 Cultural. Visiting households and firms requires permission from the relevant local district assemblies. We wrote letters of introduction and visited these assemblies to receive permissions, and this increased trust by participants. Further, we worked with the field officers to refine the design of our survey. During training activities the field officers had the opportunity to react to questions and provide suggestions for improvement. We used this feedback to make the survey as appropriate and in line with our research objectives as possible. As field officers entered the field, we received continuous feedback on ways to improve our survey and deployment procedures.

Finally, we learned that a uniform would be valuable for building trust. We provided DumsorWatch branded shirts and backpacks for the field officers so they would look official when approaching participants. These are shown in [Figure 4](#).

Figure 4: Field officers in uniform. Providing consistent branding built trust in the community as field officers visited potential participants. During the medium scale deployment, choosing a color scheme inspired by our university accidentally resulted in a color scheme similar to that of the local power utility, causing some confusion. While we were able to easily choose new colors for the large scale deployment, we highlight that it is important to consult with local experts before making branding decisions.

5.3.3 Technical. At this scale, frequently visiting sensors for debugging was no longer feasible, so we prioritized sensor stability and remote failure detection and mitigation. This included developing a full custom embedded-system for PowerWatch (shown in Figure 3 A.2) with built-in mechanisms to reset the device on failure. Additionally, we spent considerable time implementing and testing more reliable firmware, incorporating error collection libraries, and building dashboards displaying the health of both PowerWatch and DumsorWatch. We assembled this version of PowerWatch over three days with the help of a team of fellow graduate students.

Another technical challenge was dealing with mobile phone heterogeneity. We had little insight into the types of mobile phones and versions of Android among our participants. Thus, we implemented DumsorWatch to be backwards compatible to 4.0.0, a version of Android no longer supported by Google [2]. Backward compatibility took considerable engineering effort, and had side effects such as making DumsorWatch incompatible with many modern Google cloud services, including Google’s bug tracking tools.

Finally, we experienced two challenges related to SIM card operations. First, we could not identify a way to test PowerWatch sensors in the United States using the MTN postpaid SIM cards. This led to us building a United States-based testbed prior to traveling to Ghana, and performing final assembly and quality assurance in Ghana in the days leading up to the deployment. Second, MTN had not correctly provisioned the SIM cards it sold us and they subsequently could not access the network. This took multiple days of interacting MTN to fix, which delayed deployment and made clear that MTN was not well-suited to manage large fleets of SIM cards assigned to an individual customer.

These problems led us to continue exploring global SIM card options, and we tested Twilio SIM cards during this deployment. We found they had similar problems to the Particle SIMs previously evaluated. We contacted Twilio support and found their documented list of Ghanaian network operators was out of date, making unlisted providers unavailable on the Twilio network and leading to a drop in service quality.

Category	Small Scale	Medium Scale	Large Scale
Organizational	<ul style="list-style-type: none"> Local SIM procurement 	<ul style="list-style-type: none"> Hiring local staff Contracting local companies Paying outside free tier 	
Technical	<ul style="list-style-type: none"> Global SIM operation 	<ul style="list-style-type: none"> Custom hardware Firmware development App development SIM operation 	<ul style="list-style-type: none"> Assembly Site selection
Operational	<ul style="list-style-type: none"> SIM top-up 	<ul style="list-style-type: none"> Transportation Field Officers Incentivizing participants Data sharing 	<ul style="list-style-type: none"> Deployment management
Cultural	<ul style="list-style-type: none"> Learning local context 	<ul style="list-style-type: none"> Local leader approval 	<ul style="list-style-type: none"> Unexpected phone usages Survey design

Table 1: Pain points of different scales. At each scale of deployment we ran into pain points—complexities that we perceived to be more difficult than would be expected by a simple increase in deployment size. Many of these were encountered at the transition to medium scale, when local capacity needs to be built, expenses to operate the technology increase, lack of technical reliability becomes much more apparent, and systems that could once be human-operated must be automated. Large scale brings new problems, the most notable being the lack of ability to track deployment state without automated deployment management tools.

5.3.4 Operational. The operational challenges started with transporting our equipment to Ghana. We carried the PowerWatch sensors, power strips (handed out to participants as incentives), and equipment for field officers into Ghana in suitcases over multiple trips from the United States. PowerWatch sensors were carried on whenever possible to minimize any chance of being lost. This method of transportation worked, but led to multiple confrontations with airport security in the United States and customs in Ghana. We were able to overcome these hurdles by providing documentation of our project and our letters of invitation, but this transportation method depended on our team being persistent and prepared with documentation, unwrapping all equipment, labeling all equipment with tags indicating it was property of the university and not for resale, and only traveling with a few suitcases at a time.

To implement our site-selection methodology we needed GIS maps of the grid. We worked with stakeholders to determine where the best maps of the grid were maintained, and obtained these maps after repeated visits to stakeholder offices. These maps were not perfect, but included enough detail for our site-selection procedures.

At this scale we felt it was not feasible to transfer recurring incentives to participants by hand. We had anticipated this problem and had designed an incentive management system to support this goal. The system was designed to capture user behavior (e.g., whether they completed a survey, installed DumsorWatch, kept DumsorWatch installed, etc.) and transfer airtime automatically. The actual transfer of airtime took place through a third-party API. The development and testing of the incentive transfer system was done alongside deployment activities.

Figure 5: Relative locations and number of Android app events over time. Starting in August (left) we were receiving events from 989 phones in our deployment area; however, the number of participants fell to 573 by September (middle), and 310 by November (right). Because of these deployment challenges, we were unable to fully longitudinally test the app technology.

Finally, at this scale, the data collected was significant enough that stakeholders in the region began requesting access to the data. Because many of these stakeholders would be responsible for helping the project achieve further scale, we made an effort to develop and share anonymized visualizations and summary statistics.

5.3.5 Failures and Missteps. One class of failures experienced at medium scale is attributable to simple technical immaturity. For example, we found (and are still finding today) bugs both in our automated incentive transfer system and in the third-party payment API used to incentivize participants. This API is provided by a small company, but we believed it to be the best option for transferring air time in Ghana. Both technologies should have been more aggressively tested prior to launch. There is a clear need for a fleet of testing phones in Ghana against which we can implement continuous integration and automated testing of incentive transfers. However, as with most hardware-based testing systems, this is difficult to implement in practice. As a result, most participants experienced late payments, which we hypothesize caused the significant number of DumsorWatch uninstalls shown in Figure 5.

More fundamental were issues with effectively recording, connecting, and correcting critical deployment metadata. We had not anticipated the complexity of managing data about participants, devices, and app installs, each of which was collected by different systems, and some of which informed each other. This led to an ad-hoc sharing of information through our encrypted shared drive. Surveys containing participant and deployment placement information were uploaded by the field team and downloaded by the research team periodically. They were then cleaned and provided as CSV files to the individual engineer handling either sensor management or the payment system. Errors in the surveys (common due to typos in long unique IDs) were communicated back to the field team via phone calls and emails, and the resultant corrections in the field would not always be communicated back to the research team. This was ineffective while we were in Ghana and completely collapsed after we returned and could not focus full-time on deployment upkeep. As devices moved, we received multiple, conflicting reports about their current location. As a result, we permanently lost the state of some devices; five devices are still completely unaccounted for. These issues continue to make data analysis, sensor debugging, and correlation of problems with a single participant nearly impossible to manage for the devices in this deployment.

Figure 6: The PowerWatch assembly line. Over the course of four weeks, 10 undergraduates worked 110 person-hours to assemble 295 PowerWatch sensors. They were responsible for assembling the plug; screwing together the enclosure; attaching the circuit board; connecting the battery, antenna, SIM card and SD card; and provisioning the device with base firmware. They worked from team created assembly manuals and training materials.

5.4 Large-Scale Deployment

Beginning February 2019 we built upon our medium-scale deployment and added 292 new PowerWatch devices and 1,419 new app downloads in three districts of Accra, resulting in a full deployment to date of 457 PowerWatch devices and 3,400 DumsorWatch apps.

5.4.1 Organizational & Cultural. The organizational and cultural challenges did not change from the medium-scale deployment. Existing service contracts were sufficient or easily renegotiated, and the field team scaled linearly with the size of deployment.

5.4.2 Technical. The increased number and technical complexity of the new PowerWatch sensors constructed for the large-scale deployment precluded relying on other graduate students to help assemble devices as we did with the medium-scale deployment; however, the scale was still too small to be cost- or time-effective for contracted assembly. Our solution was to build our own assembly line and hire 10 undergraduates to assemble devices. This required developing training, discrete steps, and quality assurance techniques. The PowerWatch assembly line can be seen in Figure 6. Ultimately this assembly line produced the 295 PowerWatch sensors over four weeks and 110 person-hours of total work, with a 2.4% error rate, which was far below what we were anticipating. Although this activity was successful, difficulties in recruiting and paying students hourly, and challenges with the academic schedule, ensures that this model would not scale much beyond 400 units.

Similarly, the larger number of sites meant site selection was no longer easy to do by hand. This led us to develop a GIS-based site-selection system, which is able to generate sites based on our site-selection rules, label these sites, and create site location images for the field officers. This system requires the GIS maps collected from utility of the grid to be cleaned, and the system is designed and maintained by a dedicated graduate student.

We continued exploring global SIM card options, using Aeris SIM cards for a subset of this deployment. We found that due to Aeris' focus on global IoT connectivity and the number of customers they have in Sub-Saharan Africa, their SIM cards work significantly better than Particle or Twilio SIMs in Ghana.

5.4.3 Operational. The largest change was addressing operational issues described in Section 5.3.5 with custom deployment management software, described further in Section 6.1.

6 LESSONS LEARNED

Though we believe that deploying a sensor system capable of continuously collecting data was necessary to meet the goals of the deployment, it was also the greatest source of difficulty. Most of our issues can be traced back to the fact that we were deploying a continuously operating sensor network and phone application rather than just a large one-off surveying effort.

Deploying sensors and a phone application prevented us from hiring companies traditionally used to deploy large surveys, which in turn required the formation and contracting of a new company willing to manage our deployment. The design, assembly, and transportation of sensors was time-consuming and expensive. The continuous nature of the deployment generated similarly continuous data to consume, participants to incentivize, and invoices for each of these services that the university was not equipped to handle administratively. To maintain the relevance of the data produced by the sensors, we had to manage and track their status, which amplified the problems associated with errors in that collected data. We use this section to categorize these pains, provide insight into the tools we built to address them, and enumerate what we would change or expect if we were to repeat the deployment.

6.1 Continuous Monitoring Requires Continuous Upkeep

The continuous operation of a sensor network and phone application requires a significant amount of metadata and upkeep that is not required in a large survey deployment. Sensor deployment times and locations must be recorded and correlated with participant information. Unique app identifiers need to be collected to ensure app installation compliance. Participant phone numbers need to be stored so that participants can be appropriately incentivized. All of this information needs to be effectively communicated to the field officers for debugging, and updated over time because participants and their devices are in constant flux. As we describe in [Section 5.3.5](#), without a systematic approach to tracking this state at scale, the quality of the sensor deployment and our ability to properly implement our experimental design quickly degraded.

At a fundamental level, the introduction of continuous monitoring systems into a deployment introduces feedback loops that are not present in a large surveying effort. These feedback loops, shown in [Figure 7](#), have two major implications on a deployment: 1) errors introduced into the feedback loop by incorrect metadata from a survey are important, and often amplified if not addressed, and 2) state ends up existing in multiple systems, and has a high potential to become inconsistent if the feedback is not automated.

For our large deployment, we addressed each of these problems and have seen major improvements in our deployment results. The first correction was to prevent surveying errors on critical metadata. We implemented barcodes to record the unique IDs of sensors and phone applications, and we equipped the field officers with feature phones so they could text the participant to verify the participant's phone number and take a picture of the sent text message.

The second correction was custom software responsible for automatically keeping state consistent across all databases, communicating errors to the field team, and implementing corrections to survey data when submitted by the field team. In practice, the

Figure 7: The dataflow for the deployment. While traditional surveying methods have a linear data flow where data is exported for later analysis, the integration of continuous sensing in the deployment generates feedback loops which create more place in which state is stored, more necessity to communicate this state, and amplifies the issues of errors during surveying. We implement a deployment management system to alleviate these problems. Specifically, red arrows are data flows that we automate or facilitate with a deployment management tool, which we at first attempted to perform manually. Blue arrows represent data flows that we automated from the beginning because we anticipated their complexity before the medium scale deployment.

field team completes a set of deployment, retrieval, and debugging surveys in SurveyCTO [25] and the deployment management software automatically consumes these surveys using the SurveyCTO API. The data from the surveys is then verified and the information distributed to the appropriate databases. Information about surveys with errors along with a list of non-operational devices are provided to the field team through a web interface, and error corrections are communicated back to the software through a final correction survey. The deployment management software is represented by the red arrows in [Figure 7](#).

While not perfect, these techniques make the deployment significantly easier to manage. We feel that systems like these are necessary for both deploying and maintaining a continuously running sensor network, especially one in which the state of the deployment is constantly changing due to direct interaction with participants.

6.2 Global Solutions May Miss Local Context

Several times in our deployment we were forced to consider trade-offs between using technology and services developed and operated locally against similar solutions developed by larger companies targeting global scale. Specifically, we made this decision in both our choice of the cellular network provider and the service used to send realtime incentives to participants. Unsurprisingly, we found local service providers were more likely to provide high quality service

in Ghana compared to US-based companies with only nominal ability to operate globally (and little experience or market in doing so). Even our largest scale was not large enough to get dedicated support contracts with these US-based companies.

At the same time, we found local providers did not handle our medium- or large-scale deployments flawlessly. Our airtime top-up provider was not technically ready for the scale of our medium and large deployments, and neither the airtime provider nor MTN were prepared to bill and support an enterprise account. As we continue to scale we are now looking towards global-scale companies with more market share and experience in Ghana and similar geographies. We hope that these companies may provide a good mix of technical maturity, experience in handling enterprise customers, and reliability in Ghana.

6.3 University Lacks Financial Agility

One of our primary organizational problems was the inability to pay for the disparate set of services necessary to perform our deployment. This was not for lack of available funding, but due to a lack of administrative capacity in the university system.

Specifically, our university policy dictates a single day turnaround on wire transfers, while in practice this time was often over 15 days. Contracting with new companies, especially companies with which the university had never contracted before (the vast majority), often took months. The number of these contracts increased significantly because we were performing a deployment of technology that fundamentally relies on external service providers.

In practice, this meant that changes to our deployment plan—even weeks in advance—would often cause major issues. Even if we thought there was enough time for payment prior to deployment, we would still inevitably need to max out our personal ATM limits in Ghana to support deployment activities. Additionally, the university does not have good mechanisms for supporting recurring but inconsistent costs (such as a pay-per-use cloud service), because every change in cost requires approval. We found it significantly easier and more reliable to front payments for these critical services via credit card so that we could ensure they would be paid.

If we were to plan for this deployment again, we would build in significantly more time for delays and send more money than necessary to our stakeholders in Ghana early in the deployment so that they could better handle later delays in payment from the university. Still, it would be difficult to imagine the deployment running at its described pace without personal credit being extended by the research team.

6.4 Technology Usage Patterns Impact Design

Our system depends on participants to install sensors and download apps in their homes or businesses. To validate our methodology in the local context, we completed an endline survey with 462 participants from the medium scale deployment before launching the larger scale deployment. The results of this survey were surprising, and critical for the design of the next level of scale. If we were to plan this deployment again, we would run this survey earlier, as the results were important for improving system performance.

We asked participants what they thought of the sensors, the result of which is shown in Figure 8a. They liked both types of

Figure 8: Participant perception of sensors

sensors, challenging our assumption that the mobile app would seem less invasive than a physical device. Better understanding this inversion remains future work, but one hypothesis is that mobile phone resources are both scarce and highly valued.

We then explored how incentives influenced participation, the results of which are shown in Figure 8b. We find that information about reliability was valued highest. This indicates a strong local desire for energy reliability data, and suggests data alone could be effective in incentivizing participation.

Even so, many participants either uninstalled DumsorWatch from their phone or unplugged PowerWatch from the wall. We asked participants about the root cause of these behaviors. Figure 9a shows that people unplugged PowerWatch for many different reasons, some of which could likely be addressed through better information sharing (“to protect the device”, “during power outages”, “consuming too much electricity”) or through more careful user interface design (“thought it was plugged in”). These lessons will be incorporated in field officer training for future deployments.

More challenging are the results from Figure 9b, which indicate a high degree of fluidity in mobile phone usage. In particular, formatting and “flashing” (resetting) phones were significant user interactions that our team was not previously familiar with. Also, large numbers of phones broke. Our methodology never asked a participant to reinstall the app because we assumed it would stay installed, and this assumption did not map to the local context.

7 RELATED WORK

7.1 Experiences Deploying ICTDs

There is a strong tradition of research focusing on meta-insights gained from deploying information and communication technologies for development. Lessons learned as a result of these experience overlap with many of the lessons reported in this paper. These include recommendations to co-design with local practitioners or participants [9, 10, 12, 14, 23, 24], reports on the difficulty of updating, debugging, and monitoring systems with unreliable communication infrastructure [6, 8, 9, 12, 14, 20, 24], emphasis on staffing and adequately training a local team to ensure a high quality deployment [6, 9, 11, 12, 14, 24], and techniques that can be taken to ensure the sustainability of a technology [8, 11, 14, 24].

Figure 9: Why sensors were uninstalled

Our work expands on prior work by presenting our experiences and lessons learned as a function of scale, emphasizing that as scale increases, challenges related to incorrectly managing an unfamiliar local context have a higher impact on the quality of the deployment, and are harder to address post hoc.

Our deployment, where sensors were installed at participants' homes in urban and peri-urban environments, is a bit unusual in ICT for development. Many works exist within rural contexts [8, 12, 14, 15, 24], contain a dependence on user interaction [12, 14, 15, 24], or have a deployment context within larger organizations [9, 18, 28]. While our experiences do not reflect all of these other contexts, they do allow us to capture one of the first road maps for in-home, non-rural, continuous sensing in the development context.

7.2 Experiences deploying sensor networks

A 43-node sensor deployment on the uninhabited Great Duck Island was one of the first emphasizing the unattended operation of sensor systems. A key lesson was the necessity of rigorous testings as the deployment scaled from a lab environment to the field [22, 26]. The development of an easily accessible testbed described in Section 5.3 allowed us to identify a large number of errors pre-deployment and contributed strongly to our low in-field failure rate.

Anticipating problems with the communications systems critical for monitoring and debugging our sensor network, we considered employing wireless meshing to support node mobility [19, 23]. However, in practice we find that, ignoring problems described in Section 5 related to provisioning SIM cards, cellular availability is reliable enough for our application in the urban setting of Accra.

7.3 Software Systems to Support Deployments

Our experience shows software systems to track both sensor state and participant interactions are critical to support larger deployments. Other work also introduces supporting systems as first order requirements for deployment maintenance and performance. Celerate is a network management system built with an architecture and feature set “to address the real challenges we faced in building and running a Wireless ISP” [14]. The Open INcentive Kit (OINK) generalizes the development and deployment of participant incentivization systems, but many other parts of deployment management still remain to be reinvented by each practitioner as the need arises [16]. Open Data Kit (ODK) 2.0 proposes changes to the widely used Open Data Kit 1.x to move this system from a “data collection platform to a data management platform” after users reported managing the organizational complexity of implementing large survey deployments to often be overwhelming [7]. We see a potential for a similar system aiding the deployment and management of sensors and hope that the description of the deployment management system in this paper contribute to that goal.

7.4 Monitoring of the Low Voltage Network

Others have explored directly monitoring grids with small scale sensor deployments. In 1994, 20 low voltage data loggers were placed in customer residences to monitor distribution feeder systems in Buffalo, New York, USA [5], although this did not continue past the one-off deployment. More recent work proposes a sensor system for measuring low-voltage grid current and voltage at high rates [21]. While an interesting high-resolution dataset on grid performance, the cost and complexity limit wide-area deployment and are not needed to address the goals of our deployment. In contrast to these deployments, our deployment attempts measure the low voltage network at a larger scale while still maintaining independence from the utility company, and we discuss the impact of both scale and this methodology on the success of our deployment.

8 CONCLUSIONS

When first approached with the opportunity to run a deployment at scale in Accra, our team was naively confident. We were able to decompose the larger task of a deployment into subsystems, each of which we could effectively engineer. However, well-designed subsystems are not enough. Critically, we overlooked the human links between these systems, leading to problems occurring not due to the sensors malfunctioning, but instead from the complexities of sensor placement and upkeep. This meta-task of deployment management was not forgotten, but neglected for the more traditional engineering tasks, like pushing for a more fully featured firmware in PowerWatch or a better-tested implementation of DumsorWatch.

Despite this, we were able to conduct a largely successful deployment, meeting all of our design goals. This was only achieved through effort from a large and creative team, a resource that many research groups cannot easily obtain. In reaction to specific pain points, we developed meta-tools, not to replace the human links, but to assist them. We hope that with identifying and describing these tools and our broader collection of lessons learned, we have taken a step towards lowering the bar for conducting similar-scale deployments in the development community.

9 ACKNOWLEDGMENTS

This work was supported in part by the Development Impact Lab (USAID Cooperative Agreement AID-OAA-A-13-00002), part of the USAID Higher Education Solutions Network, in part by funding from the Department for International Development Energy and Economic Growth initiative (DFID EEG Project A0534A), and in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. This research has Institutional Review Board approval from the University of California at Berkeley (CPHS 2017-12-10599).

We would like to thank the anonymous reviewers for their insightful and detailed feedback and our shepherd Professor Agha Ali Raza for his guidance. This work would not have been possible without the incredible attention to detail and patience of Karen Not-sund. Many other people helped make this undertaking possible including Jeff Garnett, Hana Scheetz Freymiller, and the rest of our contacts at the Millennium Challenge Corporation; Dr. Kofi Marfo, Sharon Abena Parku, and the rest of the monitoring and evaluation team at the Millennium Development Authority (MiDA); Korshi Da Seglah and Nana Bonku Eshun at Halges; and staff from the Electricity Company of Ghana, SMEC Ghana, and ESB International. Finally, we would like to extend a special thank you to the field officers in Ghana who put in many hours working to implement and support our deployment.

REFERENCES

- [1] 2001. IEEE Guide for Electric Power Distribution Reliability Indices. *IEEE Std 1366, 2001 Edition* (2001), i–. DOI : <https://doi.org/10.1109/IEEESTD.2001.94438>
- [2] 2018. Google Play services discontinuing updates for API levels 14 and 15. (Dec 2018). <https://android-developers.googleblog.com/2018/12/google-play-services-discontinuing.html>
- [3] Ghana News Agency. 2017. President Akufo-Addo swears in new MiDA Board. <http://www.ghananewsagency.org/politics/president-akufo-addo-swears-in-new-mida-board-114271>. (2017). (Accessed on 03/15/2019).
- [4] The GridWise Alliance. 2013. Improving Grid Reliability and Resilience: Workshop Summary and Key Recommendations. <https://www.energy.gov/sites/prod/files/2015/03/f20/GridWiseImprovingElectricGridReliabilityandResilienceReportJune2013.pdf>. (2013).
- [5] Philip P Barker, James J Burke, RT Mancao, TA Short, Ch A Warren, CW Burns, and JJ Siewierski. 1994. Power quality monitoring of a distribution system. *IEEE transactions on power delivery* 9, 2 (1994), 1136–1142.
- [6] Inés Bebea-González, J Antonio Paco, Leopoldo Liñán-Benítez, F Javier Simó-Reigadas, U Rey, and J Carlos. 2011. Management framework for sustainable rural e-healthcare provision. In *Proceedings of the IADIS International Conference e-Society*. 157.
- [7] Waylon Brunette, Samuel Sudar, Mitchell Sundt, Clarice Larson, Jeffrey Beorse, and Richard Anderson. 2017. Open Data Kit 2.0: A services-based application framework for disconnected data management. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 440–452.
- [8] John Butterworth, Sam Godfrey, Ali Regah, and Scott Short. 2018. Monitoring and management of climate resilient water services in the Afar and Somali regions of Ethiopia. (2018).
- [9] Rohit Chaudhri, Darivanh Vlachos, Gaetano Borriello, Kiersten Israel-Ballard, Anna Coutoudis, Penny Reimers, and Noah Perin. 2013. Decentralized Human Milk Banking with ODK Sensors. In *Proceedings of the 3rd ACM Symposium on Computing for Development (ACM DEV '13)*. ACM, New York, NY, USA, Article 4, 10 pages. DOI : <https://doi.org/10.1145/2442882.2442887>
- [10] Salomao David, Amalia G Sabiescu, and Lorenzo Cantoni. 2013. Co-design with communities. A reflection on the literature. In *Proceedings of the 7th International Development Informatics Association Conference*. IDIA Pretoria, South Africa, 152–166.
- [11] Andy Dearden and William D Tucker. 2016. Moving ICTD research beyond bungee jumping: practical case studies and recommendations. *IEEE Technology and Society magazine* 35, 3 (2016), 36–43.
- [12] Vijay Gabale, Rupesh Mehta, Jeet Patani, K. Ramakrishnan, and Bhaskaran Raman. 2013. Deployments Made Easy: Essentials of Managing a (Rural) Wireless Mesh Network. In *Proceedings of the 3rd ACM Symposium on Computing for Development (ACM DEV '13)*. ACM, New York, NY, USA, Article 10, 10 pages. DOI : <https://doi.org/10.1145/2442882.2442896>
- [13] Paul J Gertler, Kenneth Lee, and A Mushfiq Mobarak. 2017. Electricity Reliability and Economic Development in Cities: A Microeconomic Perspective. (2017).
- [14] Shaddi Hasan, Yahel Ben-David, Max Bittman, and Barath Raghavan. 2015. The Challenges of Scaling WISPs. In *Proceedings of the 2015 Annual Symposium on Computing for Development (DEV '15)*. ACM, New York, NY, USA, 3–11. DOI : <https://doi.org/10.1145/2830629.2830637>
- [15] Kurtis Heimerl, Janani Vasudev, Kelly G Buchanan, Tapan Parikh, and Eric Brewer. 2010. A case study on designing interfaces for multiple users in developing regions. In *Proceedings of the First ACM Symposium on Computing for Development*. ACM, 9.
- [16] Noah Klugman, Santiago Correa, Pat Pannuto, Matthew Podolsky, Jay Taneja, and Prabal Dutta. 2019. The open incentive kit (OINK): standardizing the generation, comparison, and deployment of incentive systems. In *Proceedings of the Tenth International Conference on Information and Communication Technologies and Development*. ACM, 17.
- [17] Noah Klugman, Javier Rosa, Pat Pannuto, Matthew Podolsky, William Huang, and Prabal Dutta. 2014. Grid watch: Mapping blackouts with smart phones. In *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications*. ACM, 1.
- [18] Zahir Koradia, C. Balachandran, Kapil Dadheech, Mayank Shivam, and Aaditeswar Seth. 2012. Experiences of Deploying and Commercializing a Community Radio Automation System in India. In *Proceedings of the 2Nd ACM Symposium on Computing for Development (ACM DEV '12)*. ACM, New York, NY, USA, Article 8, 10 pages. DOI : <https://doi.org/10.1145/2160601.2160612>
- [19] Ting Liu, Christopher M Sadler, Pei Zhang, and Margaret Martonosi. 2004. Implementing software on resource-constrained mobile sensors: Experiences with Impala and ZebraNet. In *Proceedings of the 2nd international conference on Mobile systems, applications, and services*. ACM, 256–269.
- [20] Yunhao Liu, Yuan He, Mo Li, Jiliang Wang, Kebin Liu, and Xiangyang Li. 2013. Does wireless sensor network scale? A measurement study on GreenOrbs. *IEEE Transactions on Parallel and Distributed Systems* 24, 10 (2013), 1983–1993.
- [21] Heiko Maaß, Hüseyin Kemal Cakmak, Felix Bach, Ralf Mikut, Aymen Harrabi, Wolfgang Süß, Wilfried Jakob, Karl-Uwe Stucky, Uwe G Kühnapfel, and Veit Hagenmeyer. 2015. Data processing of high-rate low-voltage distribution grid recordings for smart grid monitoring and analysis. *EURASIP Journal on Advances in Signal Processing* 2015, 1 (2015), 14.
- [22] Alan Mainwaring, David Culler, Joseph Polastre, Robert Szewczyk, and John Anderson. 2002. Wireless sensor networks for habitat monitoring. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*. ACM, 88–97.
- [23] Nithya Ramanathan, Laura Balzano, Deborah Estrin, Mark Hansen, Thomas Harmon, Jenny Jay, William Kaiser, and Gaurav Sukhatme. 2006. Designing wireless sensor networks as a shared resource for sustainable development. In *2006 International Conference on Information and Communication Technologies and Development*. IEEE, 256–265.
- [24] Carlos Rey-Moreno, Zukile Roro, William D. Tucker, Masbulele Jay Siya, Nicola J. Bidwell, and Javier Simo-Reigadas. 2013. Experiences, Challenges and Lessons from Rolling out a Rural WiFi Mesh Network. In *Proceedings of the 3rd ACM Symposium on Computing for Development (ACM DEV '13)*. ACM, New York, NY, USA, Article 11, 10 pages. DOI : <https://doi.org/10.1145/2442882.2442897>
- [25] SurveyCTO. 2019. Because your data is worth it | SurveyCTO. <https://www.surveyccto.com/>. (2019). (Accessed on 03/15/2019).
- [26] Robert Szewczyk, Joseph Polastre, Alan Mainwaring, and David Culler. 2004. Lessons from a sensor network expedition. In *European Workshop on Wireless Sensor Networks*. Springer, 307–322.
- [27] Jay Taneja. 2017. *Measuring Electricity Reliability in Kenya*. Technical Report. Working paper, available at <http://blogs.umass.edu/jtaneja/files/2017/05/outages.pdf>.
- [28] Heather Underwood, S. Revi Sterling, and John K. Bennett. 2013. The Design and Implementation of the PartoPen Maternal Health Monitoring System. In *Proceedings of the 3rd ACM Symposium on Computing for Development (ACM DEV '13)*. ACM, New York, NY, USA, Article 8, 10 pages. DOI : <https://doi.org/10.1145/2442882.2442893>