

Chapter 6

Measuring Grid Reliability in Ghana



Noah Klugman, Joshua Adkins, Susanna Berkouwer, Kwame Abrokwah, Matthew Podolsky, Pat Pannuto, Catherine Wolfram, Jay Taneja, and Prabal Dutta

1 The Development Challenge

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

N. Klugman (✉) · J. Adkins · P. Dutta
Electrical Engineering and Computer Science, University of California, CA, USA

nLine, Inc, Berkeley, CA, USA
e-mail: nklugman@berkeley.edu; adkins@berkeley.edu; prabal@berkeley.edu,
<https://www.nline.io>

S. Berkouwer
The Wharton School, University of Pennsylvania, Philadelphia, PA, USA
e-mail: sberkou@wharton.upenn.edu

K. Abrokwah
Brixels Company Limited, Accra, Ghana

M. Podolsky
Electrical Engineering and Computer Science, University of California, CA, USA

P. Pannuto
Computer Science and Engineering, University of California, San Diego, CA, USA
e-mail: ppannuto@ucsd.edu

C. Wolfram
Haas School of Business, University of California, Berkeley, CA, USA
e-mail: cwolfram@berkeley.edu

J. Taneja
Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, USA
e-mail: jtaneja@umass.edu

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 have often focused on increasing access to electricity by expanding the grid. There has been less focus on increasing the quality of electricity provided by the existing grid. However, poor reliability is often associated with a reduction in the demand, utilization, and social benefit of electricity (Gertler et al., 2017). Reassuringly, the importance of grid reliability is increasingly being recognized: the UN Sustainable Goals now specify that access to electricity must also be reliable (McCollum et al., 2017).

Improvements in electricity reliability can be harder to achieve than improvements in access. Improving a grid's reliability requires fine-grained information about how different attributes of the grid perform. However, these data are historically expensive to collect and not prioritized (The GridWise Alliance, 2013; Taneja, 2017). Many electrical utilities, investors, and energy regulators are under-resourced for the enormously complex and expensive task of planning, extending, and operating the grid at consistent reliability of service, especially in the presence of rising consumption and increasingly adverse weather patterns. Thus, the development challenge is to reduce the economic and technical barriers to monitoring the quality of existing electricity networks in the developing world, so that this data can inform actions taken to improve grid reliability.

Developing a novel technology in an academic laboratory setting is not sufficient for addressing critical information gaps: The widespread deployment of these technologies brings a unique and unpredictable set of challenges. This chapter describes a project that transferred research out of the academic laboratory and into a field setting. Our interdisciplinary team of engineers and economists designed, deployed, and continues to operate a large sensor network in Accra, Ghana, that measures power outages and quality at households and firms. The deployment began in June 2018 and consisted of 457 custom sensors, nearly 3500 mobile app installations, and nearly 4000 participant surveys, as well as user-incentive and deployment management meta-systems developed by our team, which will continue collecting reliability data until at least September 2021.¹ Grid reliability data collected by our deployment in Accra will enable multiple impact evaluations of investments in the grid that, in part, explore the effects of improvements in reliability on socioeconomic well-being.

Entering the field, we anticipated we would have to pivot our technology and deployment methodology, both initially designed in the lab, to respond to unanticipated cultural, technological, and organizational requirements. This indeed happened, as we describe throughout this case study. However, we also held the incorrect assumption that the non-technology dimensions of this work would stabilize over time and wouldn't meaningfully contribute as a barrier to scale as compared to overcoming engineering hurdles. However, in practice, we found that

¹ This research has Institutional Review Board approval from the University of California at Berkeley (CPHS 2017-12-10599).

pain points experienced at each deployment scale were fundamentally interdisciplinary, often manifesting as new or more stringent cultural and organizational requirements, each of which required innovation to overcome. This case study therefore focuses mainly on the consequences of changing deployment context, in this case by changing scale, and distilling lessons we wish we had learned before forecasting budgets, human resource requirements, and project timelines. These lessons remain critical for our team as we continue to scale our deployment beyond the work described here. We hope that learning about our case study encourages other researchers and helps them hit the ground running as they pursue highly granular measurements of the world's critical systems.

2 Context

We briefly introduce the context in our deployment site of Accra, Ghana, touching on the local energy context and introducing our implementing partners. We then introduce the research questions that data from our deployment was designed to explore, providing further context for the design decisions described in the rest of this case study.

2.1 *Energy Environment in Ghana*

Ghana's electric grid has roughly 4.74 million connections and experiences a peak load of 2881 MW, a supply capacity of 4695 MW, and an estimated 24.7% distribution loss rate (Energy Commission Ghana, 2019). The distribution utility in the capital city of Accra is the Electricity Company of Ghana (ECG).

Reliable electricity has the potential to provide huge social and economic benefits (Bowers, 1988; McCollum et al., 2017; Hamidu & Sarbah, 2016). In Ghana, however, the grid at times falls short of enabling these, causing frustration that has even boiled over to civil unrest (Aidoo & Briggs, 2019; Ackah, 2015). From 2013 through 2015, the country experienced drastic electricity undersupply, culminating in outages between 6 and 24 h during 159 days in 2015. This period is known as "Dumsor," the Twi word for "off-on." While Dumsor has largely been remedied with the introduction of new generation capacity (Clerici et al., 2016; Millennium Challenge Corporation, 2008), Ghana still reports electricity reliability metrics that underperform relative to countries with similar GDPs (Millennium Development Authority, 2018).

2.2 Partners

Mirroring the increasingly global focus on energy reliability, current investments in the Ghanaian electricity sector are primarily aimed at improving the *reliability* of electricity distributed on the existing grid. Partially in response to the Dumsor crisis, the Millennium Challenge Corporation (MCC) and the Government of Ghana signed the Ghana Power Compact in 2014, a USD 308 million² investment designed to improve the grid generation, transmission, and distribution systems in Ghana, to be implemented by the newly created Millennium Development Authority (MiDA) (Millennium Development Authority, 2018). This investment has multiple goals, including cutting operational costs, reducing transmission and distribution losses, increasing affordable access to grid connections, and improving reliability. The country's current work to improve grid reliability motivated our selection of Ghana as our deployment site (Millennium Challenge Corporation, 2008). One of the goals of our deployment is to work with both MCC and MiDA to provide data to support their monitoring and evaluation goals.

2.3 High-Level Goals

2.3.1 Improving Energy-Reliability Data Quality

$$\text{SAIDI} = \frac{\text{Total duration of sustained interruptions in unit time}}{\text{Total number of consumers}} \quad (6.1)$$

$$\text{SAIFI} = \frac{\text{Total number of sustained interruptions in unit time}}{\text{Total number of consumer}} \quad (6.2)$$

To improve reliability, it is important to measure it (McCollum et al., 2017; Sustainable Energy for All, 2018). Two widely used key performance indicators that capture the overall reliability of the grid are the system average interruption duration index (SAIDI) and the system average interruption frequency index (SAIFI) (IEEE Std 1366, 2001). The calculation of SAIDI and SAIFI is shown as Eqs. (6.1) and (6.2). The monitoring and evaluation teams at MCC and MiDA consider changes in SAIDI and SAIFI to be indicative of the impact of the Ghana Power Compact (Millennium Development Authority, 2018).

² When the Compact was first signed in 2014, it totalled USD 498 million – this amount was lowered to USD 308 million in 2019 when a 20 year concession of the utility was cancelled.

In Accra, the highest spatial and temporal resolution measurements of grid interruptions are collected by the supervisory control and data acquisition (SCADA) system operated by ECG. This SCADA system has limited reach, covering only high-voltage transmission lines and some portion of the medium-voltage distribution network (Nunoo & Ofei, 2010). To improve monitoring, ECG has recently deployed some smart meters, but economic and social challenges raise barriers to achieving broad smart meter coverage in the short term (Acakpovi et al., 2019; Millennium Development Authority, 2014; Banuenumah et al., 2017). ECG recently completed a much larger effort to upgrade their prepaid meters; however, many of the prepaid meters do not regularly communicate power quality measurements back to the utility (Quayson-Dadzie, 2012; Electricity Company of Ghana., 2019).

Measurements of low-voltage outages come primarily from customer calls, yet data collected from the national call center suggest that this data stream is imperfect. Frequent outages may reduce willingness to report, as calling the utility to report an outage can require significant effort without a guarantee that it will shorten the duration of an outage: 90% of respondents to a survey we conducted had not reported a single outage to ECG in the preceding 3 months, despite also acknowledging in the same survey that they had experienced outages during that time. There are few call center reports during the night while people are asleep, and dips in reporting during the day while people are at work. While some of these patterns may reflect reality (a grid will fail more often when it is hot and being used at capacity, conditions more likely in the middle of the day (He et al., 2007)), these data are likely not truly representative.

Therefore, our deployment aims to generate more accurate estimates of SAIDI and SAIFI by placing sensors in the field that extend beyond the SCADA system to automatically report the location and duration of power outages in the low-voltage network, as well as further upstream. We also capture voltage fluctuations and frequency instabilities at the low-voltage level, as these can have significant impacts on the value of appliances and machinery.

2.3.2 Providing Utility-Independent Measurements

To understand how well infrastructure improvements impact reliability, it is important to have utility-independent measurements (Millennium Development Authority, 2018). Many widely used tools, including SCADA and smart meter technologies, are dependent on utility participation, in part because they directly interface with utility property. From an academic perspective, independence is important as it allows for unbiased research output. Independence is often desired by regulators as well, who may want to verify measurements provided by the utility, as the utility has incentives to report favorable reliability metrics. Currently, no high-resolution source of independent data about grid reliability in Accra exists. Our deployment was designed to evaluate the feasibility and efficacy of a novel sensing methodology for monitoring the reliability of the electricity grid while working independently of the utility. Our physical sensor is designed to be installed at outlets in households

and business, and our sensing app personal smartphones, allowing us to choose deployment sites and deploy sensors without utility involvement. The data returned from our deployment is truly independent.

2.3.3 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, such as fans and refrigerators, or income-generating assets, such as welding, sewing, or other productive machinery. Our deployment was designed in part to generate both reliability and socioeconomic data to support an ongoing economic study that aims to estimate the causal impact of power quality and reliability on socioeconomic outcomes, such as productivity and health, for residents and firms in Accra.

3 Innovate, Implement, Evaluate, Adapt

Starting in May 2018, we conducted three deployments at three different scales: a small-scale pilot, a medium-scale deployment, and a large-scale deployment.³ To date, 3400 individuals in Accra downloaded our mobile app, called DumsorWatch; 457 people installed our plug-in sensor, called PowerWatch; and over 4000 surveys were performed to directly measure socioeconomic outcomes. In December 2019, we surveyed an additional 462 participants to understand their experiences and to collect updated measures of time-varying socioeconomic outcomes.

We present the technology deployed, the design of our deployment, and where our planning and assumptions failed or caused unexpected problems. While doing so, we attempt to categorize our key challenges and describe the steps we have taken to overcome each of these challenges, emphasizing issues that arose that were more complex and/or costly than originally forecasted. We find that our experiences differed depending on the scale of the deployment, each scale uncovering its own complexities.

3.1 Innovation: Data Collection Instruments

We developed two types of data collection instruments – sensors and surveys – to achieve the goals described in Sect. 2.3. These instruments collect the required data to estimate SAIDI and SAIFI independently of utility participation.

³ As of November 2020, we continue to expand with additional deployments.

3.1.1 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 sensor called PowerWatch that is plugged into a power outlet at a household or business.

DumsorWatch is an Android app installed on the everyday-use smartphone of a participant who lives and/or works in Accra. “Dumsor,” the local word for power outages, was used for branding and association with power outages in the Ghanaian context. DumsorWatch automatically senses power outages and power restorations through a combination of on-phone sensors and cloud services (Klugman et al., 2014).

PowerWatch, our plug-in sensing technology, integrates power reliability sensors with a GSM radio to send measurements in near-real time to a cloud database (Fig. 6.3). By designing PowerWatch to plug into a participant’s home or business, as opposed to connecting directly to the electric grid, we avoid the need for prior approval or cooperation to deploy the sensors and therefore maintain independence from the utility, a primary goal of our deployment. PowerWatch senses power outages and power restorations timestamped to the millisecond, as well as GPS-based location, voltage, and grid frequency.⁴ PowerWatch contains a battery to allow for continuous reporting throughout a power outage and will queue data if there are GSM-network-connectivity problems, to be uploaded once GSM connectivity is restored.

How Do PowerWatch and DumsorWatch Work?

PowerWatch consists of an outage detection sensor that plugs into an outlet and is installed in homes and businesses and reports the state of the grid over a cellular backhaul to the cloud. Every minute, the sensor takes a reading of power state, grid voltage, grid frequency, GPS, and cellular quality. It also records the number of nearby Wi-Fi signals as secondary validation, as wireless hotspots may be grid powered. In addition, upon changes in power state, the device records the timestamp (from an on-board real-time clock) and current acceleration. All these measurements are stored locally on an SD card and transmitted to the cloud when a cellular connection is available. Acceleration signals that a participant is interacting with the device, making it likely that any charge-state change at that time is a false positive and allowing us to more easily reject the data point. The sensor contains a 2000 mAh battery, which can run the sensor for several days, longer than most outages

(continued)

⁴ PowerWatch internally records voltage measurements 150 times per AC cycle; then every 2 min, it reports the RMS voltage and grid frequency averaged over 30 AC cycles.

in Accra. When the sensor is on battery power, it still reports data at the same frequency to our servers, a feature necessary for calculating outage duration.

The primary sensors used by the DumsorWatch app are the phone's location sensors (GPS), the phone's charging state which monitors if the phone is connected to a power source or not (and notifies DumsorWatch when that state changes), and the system clock to give the time of any observed events. Secondary sensors in the app help refine the likelihood that a change in charge state corresponds to a power outage or restoration. For example, the accelerometer can measure if the phone was moving when its charge state changed (as it would if a charging cable was inserted or removed from the phone side), and the phone's Wi-Fi radio can report the presence or absence of wireless hotspots. On-phone processing and analysis of microphone recordings may be able to detect the presence of a 50 Hz "hum" of grid mains. Additionally, the users of the app can manually report a power outage and power restoration by pressing a button in the app.

For both technologies, a cloud-based analytic system searches for outage reports from multiple devices to ensure the validity of an outage. To perform this search, we cluster outage reports into density-based clusters in both space and time. This lets us reject noise from a single sensor (i.e., a single participant unplugging a device or a pre-paid meter running out of credit) and ensures that only true outages are reported by the system.

3.1.2 Surveys

A 60-min socioeconomic survey accompanied the deployment of each PowerWatch device, and all participants who received a PowerWatch device also downloaded the DumsorWatch app. A shorter survey was administered to participants who solely downloaded the DumsorWatch app. All surveys were completed using SurveyCTO, and all participants received an airtime transfer as a thank you for participation (SurveyCTO, 2019). We conducted 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 for the economic study, the survey was used to support the development and deployment of the technology itself. For example, the survey recorded a unique code for the PowerWatch device and DumsorWatch app and the participant's phone number and GPS location. To inform DumsorWatch debugging,

we asked how they used their mobile phones, how many phones and SIM-cards they use, and how frequently they upgrade their phones. To inform the PowerWatch deployment, we asked whether the participant turns off their electricity mains at night and whether they had any safety concerns about PowerWatch.

3.2 Innovation: Deployment Methodology

To support our deployment as scale increased, we designed and implemented a novel set of deployment management tools. While our methodology evolved to support each deployment scale, its general structure remained fairly consistent. First, we developed criteria for site selection that allows us to answer specific socioeconomic questions. Next, we devised a sampling procedure that gave sufficient coverage of each chosen site, as well as sufficient redundancy to enable cross-validation of the new measurement technology. Finally, we worked with a team of field officers to deploy in the chosen sites, employing our deployment management tools to maintain and monitor the system. The rest of this section considers each of these components in detail.

3.2.1 Site Selection

We selected a subset of the sites where infrastructure upgrades are planned (“treatment sites”) and then quasi-randomly selected a set of sites that are comparable in observable characteristics (“control sites”). For each site, we defined a geographic surveying area that is the intersection of a 200-meter radius from the site centroid and a 25-meter region extending from the low-voltage network being measured. We wanted the area to be relatively small, so that we could have a high degree of confidence that customers within the area were all connected to the same infrastructure, but it needed to be large enough to have a sufficient number of residents or firm owners for us to enroll into the study. We performed this analysis using GIS tools operating on a newly created map of Accra’s grid constructed by an independent firm that had been contracted by MiDA as part of the improvements funded by the Ghana Power Compact. Using these GIS maps, we produce a series of maps marking the geographic area bounding each site. Field officers used these maps, along with the GPS coordinates for the sites, to identify the surveying area and deploy sensors accordingly.

3.2.2 Sampling Strategy

We deployed our sensors at the home or place of work (or both, if these are co-located) of Accra residents, targeting a 50/50 split between households and firms. Installing PowerWatch at consumer plugs and DumsorWatch on consumer phones

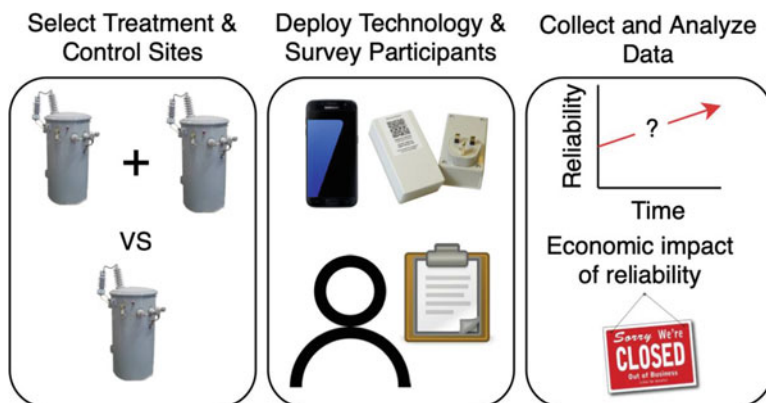


Fig. 6.1 Deployment overview

Notes: 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

allows us to not depend on direct access to utility infrastructure such as transformers or lines and to measure power quality at the point where it is least understood: the customer (Fig. 6.1).

We planned a deployment of three PowerWatch devices and 20 DumsorWatch app downloads at each site. Our strategy is built around redundant sampling such that multiple sensors are placed under a single transformer. When multiple 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 transformers reporting outages simultaneously, we can infer the outage occurred at a higher level of the grid. This sampling strategy is shown in Fig. 6.2.

3.2.3 Deployment and Surveying Team

We hired local staff who supported our continuously operating deployment. One team member works full time as a field manager to oversee initial roll-out and ongoing maintenance of the system and an auditor to follow up with participants who report problems or whose sensors are no longer functioning.

To implement our medium- and large-scale deployments, we temporarily employed a team of ten field officers and three team leads. Prior to the start of the deployment, the field officers were trained extensively to ensure the correct protocols were used to obtain consent, conduct surveys, plug in the power strip and PowerWatch device at the respondent's home, download the app onto the respondent's phone, and conduct any necessary troubleshooting related to the

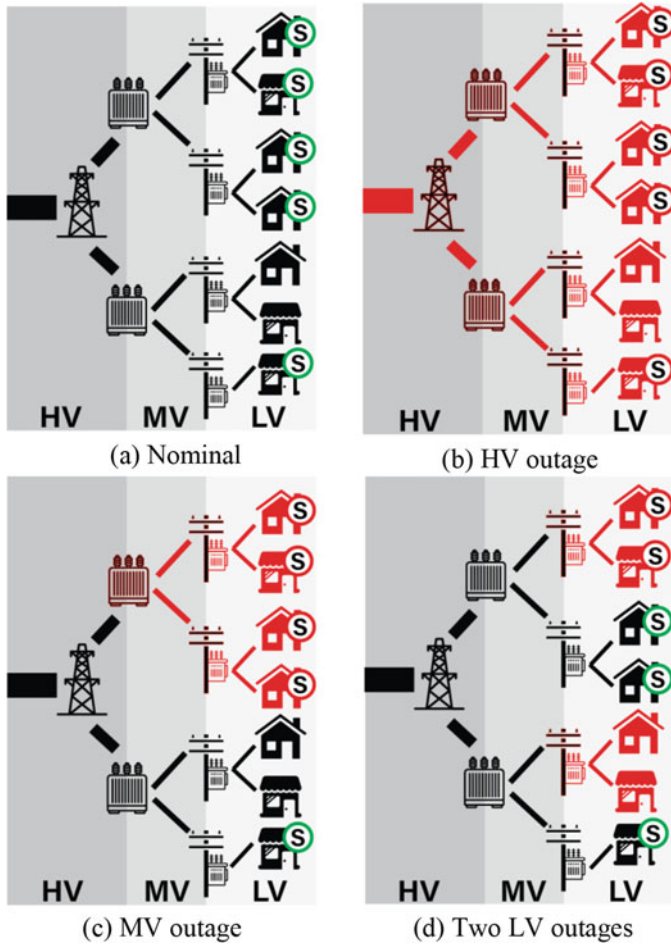


Fig. 6.2 Deployment methodology for sensors

Notes: Sensors, either the PowerWatch sensor or DumsorWatch app (both labeled S), monitor power at homes and firms. By randomly sampling outlets at households and businesses, they detect a significant portion of low-voltage (b), medium-voltage, (c) and high-voltage (d) outages. In the aggregate this data can be used to estimate the average frequency and duration of outages, including both single-phase and multi-phase outages, by looking at reports from sensors that are close together in space and time. Additionally, this spatiotemporal analysis allows identification of what voltage level experienced the fault. Undersampling can lead to missed outages when sensors are not present in any of the affected units, shown in the bottom outage of (d)

technologies. Field officers find potential participants, get informed consent, and screen their eligibility. They then conduct the survey, install the sensors, and answer any participant questions. We conducted multiple training exercises where each team member learned about the technologies being deployed and practiced completing the survey and deploying the technologies.




Fig. 6.3 Evolution of PowerWatch with each deployment

Notes: 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



Fig. 6.4 Field officers in uniform

Notes: 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 easily able to choose new colors for the large-scale deployment, we highlight that it is important to consult with local experts before making branding decisions

Field officers visited sites in groups of two to alleviate safety concerns. We provided team uniforms, shown in Fig. 6.4, to make clear they are part of an official project. We also provided backpacks to carry supplies, tablets to conduct the survey, Wi-Fi 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.

3.2.4 Dependence on Participants

The placement of PowerWatch sensors directly in homes and firms – where participants may unplug them, run generators, or experience power shutoffs due to nonpayment – increases the noise of our data relative to a deployment on utility-owned equipment such as transformers. Similarly, the DumsorWatch app may be uninstalled from respondents' phones, reducing coverage and leading to a potentially under-sampled signal. A key challenge was thus ensuring that we only enrolled participants who had the ability and desire to participate for the full study duration and then to minimize any cause for participants to choose to withdraw consent from participation.

In a preemptive attempt to decrease the statistical noise caused by human factors, we screen participants for specific criteria including owning a phone with Android version 4.1–8.1 and being an active customer on the grid. In order to minimize attrition, we explain the goals, risks, and benefits of the project, as part of the consent process. Finally, we provide a phone number to call if participants have any questions or concerns.

To further encourage continued participation, we compensate participants monthly with airtime credits on their mobile phone. All participants whom we recruited to download the DumsorWatch app received 5 Ghana Cedi (0.93 USD) of airtime for initial recruitment and 4 Ghana Cedi (0.75 USD) monthly for keeping DumsorWatch installed. Participants who also installed a PowerWatch device received an additional 10 Ghana Cedi (1.86 USD) for installing the sensor and 5 Ghana Cedi (0.93 USD) monthly for keeping PowerWatch installed. 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.

3.2.5 Deployment Management Tools

We developed three software subsystems to support the deployment: (1) an automated incentive system to transfer the airtime incentives; (2) a deployment management system to (a) track sensor and participant status and (b) display deployment health to the field management team; and (3) a data visualization and analysis system. We discuss these systems, and the experiences that led us to develop them as the deployment scaled, in Sect. 3.5.

3.3 *Evaluation: Overview*

For each of the small-, medium-, and large- scale deployments, we report problems that occurred, techniques used to mitigate their impacts, and the effectiveness of the mitigation. To emphasize the parallels between challenges exposed at

different scales, we organize this discussion around four categories of challenges: organizational, cultural, technical, and operational. Organizational challenges relate to procurement, hiring, and finances; cultural challenges relate to how cultural considerations impacted the deployment and operation of the technology; technical challenges relate to the development, manufacturing, and functioning of the technology; and operational challenges relate to the successful deployment and operation of the technology.

3.4 Evaluation: Small-Scale Pilot

The first activity we performed was a deployment of 15 PowerWatch sensors and 5 DumsorWatch app downloads. 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: devices were not deployed with participants enrolled from the public but in the private homes of our research partners. The primary challenges were related to producing the technology, connecting the PowerWatch sensors to the cellular network, and building enough local capacity to deploy PowerWatch and DumsorWatch.

In addition to testing the technology, we worked to build relationships to support future scaling. We reached out to local stakeholders for feedback on the assumptions driving our sensor design, speaking with engineers and managers at ECG, MiDA, and several independent contractors involved in the Ghana Power Compact. We also received data from ECG that helped validate our hypothesis that their existing estimates of SAIDI and SAIFI could benefit from higher-resolution measurements.

Even at a small scale, we experienced unanticipated technical challenges. To connect the PowerWatch devices to the cellular network, we initially used SIM cards sold by Particle, the US-based manufacturer of the cellular technology used in PowerWatch, in part because these SIM cards had been advertised to work in any country. But in practice, their ability to maintain a stable network connection was worse than that of a local SIM. We therefore decided to use SIM cards from the largest local carrier (MTN), but we encountered a three-SIM-card-per-person limit upon purchase. Although we were able to circumvent this by visiting different stores, purchasing SIM cards in stores was not an option for future scale.

Another challenge was keeping SIM cards functional over the full study period. Prepaid SIM cards require data plans, which are purchased using an unstructured supplementary service data (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.

Table 6.1 Pain points of different scales

Type	Small scale	Medium scale	Large scale
Organizational	Local SIM procurement	Hiring local staff	
		Contracting local companies	
		Paying outside free tier	
Technical	Global SIM operation	Custom hardware	Assembly
		Firmware development	Site selection
		App development	
		SIM operation	
Operational	SIM top-up	Transportation	Deployment management
		Field officers	
		Incentivizing participants	
		Data sharing	
Cultural	Learning local context	Local leader approval	Unexpected phone usages
			Survey design

Notes: At each deployment scale, we encountered pain points – complexities that were more difficult than we expected from a simple increase in deployment size. Many occurred at the transition to medium scale, when local capacity had to be built, expenses to operate the technology increased, lack of technical reliability became more apparent, and systems that could once be human-operated had to be automated. Large scale brought new problems, most notably the need for automated deployment management tools to track deployment state

3.5 *Adaptation from Small-Scale Pilot Experience and Evaluation of Medium-Scale Deployment*

In our medium-scale deployment, 1981 individuals downloaded the DumsorWatch app, and 165 individuals installed PowerWatch sensors. After an initial 1-week training, field officers first deployed the PowerWatch sensors, which included also administering a detailed socioeconomic survey and installing the DumsorWatch apps on the phones of PowerWatch recipients. Once the deployment of the PowerWatch sensors was complete, field officers spent 3 weeks deploying the DumsorWatch apps among an additional set of participants, which also included a shorter socioeconomic survey. Once initial deployment was complete, the monitoring activity continued for 7 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 implementing the survey instruments. We enumerate the changes experienced as we increased from small- to medium-scale in Table 6.1, paying particular attention to the challenges extracted.

3.5.1 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 effort from the research team. Bureaucratic requirements at the university also caused frequent delays in payments, especially when payment was needed in a short time frame (1–2 weeks). Increased flexibility, for example, at an independent organization or private sector company, might better suit technological deployments that have some complexity in administrative streams like finances and hiring.

Because prepaid SIM cards were not available at the quantities we now needed, we had to enter into a contract with the cellular provider, MTN. To alleviate concerns about whether our application was legitimate, we visited the MTN main office in our university shirts, gave a technical demo, and answered questions about our backgrounds and affiliations.

At medium scale, many of the cloud-based software services our systems were built on were no longer eligible for free-tier usage. For one service, this meant that we would be unable to continue without signing a multiyear 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.

3.5.2 Cultural

Visiting households and firms requires permission from the relevant local district assemblies. We wrote letters of introduction and visited these assemblies to receive permission. Receiving this permission also increased participant trust.

We also worked with the field officers to refine our survey design. 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 culturally 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 participant trust. We provided DumsorWatch-branded shirts and backpacks for the field officers, so they would look official when approaching participants. These are shown in Fig. 6.4. Field officers also carried identification cards that they could show participants in case of any questions.

3.5.3 Technical

At medium 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 Fig. 6.3 B.1) 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 3 days with the help of fellow graduate students.

Another technical challenge concerned mobile phone heterogeneity. We had little insight into what types of mobile phones and versions of Android were most common in Accra. Thus, we implemented DumsorWatch to be backwards compatible to 4.0.0, a version of Android no longer supported by Google (Google, 2018). 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, making app failures much harder to correct. Further, we chose to support older versions of Android at the expense of supporting the newest Android version at the time, Android 8.1, a decision that made those with the newest phones no longer eligible for participation. While this did not reject large numbers of participants during this deployment, as more people moved to newer devices, it would have more impact on recruitment, necessitating future engineering costs.

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 MTN postpaid SIM cards, which were not configured to allow PowerWatch to connect to cellular networks outside of Ghana. We therefore built a US-based testbed for sensor development that used US-based SIM Cards. However, to perform final assembly and quality assurance, steps that require PowerWatch to be connected to the MTN network, we needed to wait until the sensors were in Ghana for a deployment, compressing the timeline of these tasks and increasing risk if quality assurance failed. Second, MTN's process for provisioning the SIM cards required additional oversight and took much longer than expected, which delayed deployment and made clear that a different partner would be required to manage large fleets of SIM cards assigned to a single customer.

These problems led us to continue exploring global SIM card options, and we tested a small number of 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. This theme of global solutions lacking local service quality is explored further in Sect. 4.2.

3.5.4 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 the plane whenever possible to minimize their chance of being lost. This method of transportation worked but led to multiple questions from airport security in the United States and customs in Ghana. We were able to overcome these hurdles by creating documentation about our project and providing this along with letters of invitation from MiDA, but even still this transportation method depended on our team being persistent and prepared with documentation, unwrapping all equipment from its packaging to make it look less likely to be resold, 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. More generally, ensuring safe and timely transport of nonconsumer technology across borders will likely require additional measures depending on the local context.

Implementation of our site selection methodology required GIS maps of the electric grid. We worked with stakeholders to determine where the best maps of the grid were maintained, a task made more complicated by maps being held by multiple subcontractors of the Ghana Power Compact. With MiDA's support we were given access to maps that, while not perfect, included enough detail for our site selection procedures. At this medium scale, which was also relatively concentrated geographically, it was feasible for a member of the research team to study the GIS maps visually, identify the proposed treatment sites, manually identify control sites quasi-randomly to match treatment sites based on observable characteristics such as grid characteristics and satellite imagery, and produce site maps that field officers could use to identify potential respondents at each site. This process met the requirements for this scale deployment but would prove exceedingly complicated for a larger deployment.

At medium scale we felt it was not feasible to transfer recurring incentives to participants by hand. We had anticipated this problem and designed an incentive-management system to support this goal. The system was designed to capture user behavior (e.g., whether a participant completed a survey, installed DumsorWatch, kept DumsorWatch installed, etc.) and to transfer airtime automatically. The actual transfer of airtime took place through a third-party API. We developed and tested the incentive transfer system alongside our deployment activities (Klugman et al., 2019).

Finally, at medium scale, the data collected were significant enough that they became valuable to stakeholders in the region. 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.




Fig. 6.5 Relative locations and number of Android app events over time

Notes: Starting in August 2018 (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

3.5.5 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 airtime 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 for 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 Fig. 6.5.

More fundamental were issues with effectively recording, connecting, and correcting critical deployment metadata, such as the location and status of each device and payment, which we collectively refer to as the *state* of the deployment. We had not anticipated the complexity of managing data about participants, devices, and app installs, each of which was collected by a different system and some of which informed each other.

This led to an ad hoc sharing of information through our encrypted shared drive. The field team uploaded surveys containing participant- and deployment-placement information on a daily basis. The research team downloaded and cleaned these periodically and provided the resulting 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 were not always communicated back to the research team. This process was ineffective while we

were in Ghana and completely collapsed after we returned to the United States 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 unaccounted for. These issues continue to make data analysis, sensor debugging, and correlation of problems with specific participants difficult for the devices in this deployment.

3.6 Adaptation from Medium-Scale Deployment and Evaluation of Large-Scale Deployment

Beginning in February 2019, we built upon our medium-scale deployment and added 292 new PowerWatch devices and 1419 new app downloads in 3 districts of Accra, resulting in a combined large-scale deployment of 457 PowerWatch devices and 3400 DumsorWatch apps.

3.6.1 Organizational and 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.

3.6.2 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 us to develop discrete assembly steps, a training guide, and quality assurance techniques. The PowerWatch assembly line can be seen in Fig. 6.6. This assembly line produced the 295 PowerWatch sensors over 4 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, mean this model would not scale much beyond 400 units.

The larger number of sites meant site selection was no longer easy to do manually. This led us to develop a GIS-based site selection system, which generates sites based on our site selection rules, labels these sites, and creates site location images for the field officers. This system requires cleaning the GIS maps of the grid collected from the utility and was designed and maintained by a dedicated graduate student.



Fig. 6.6 PowerWatch assembly line

Notes: Over 4 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

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 in Ghana than Particle or Twilio SIMs.

3.6.3 Operational

The largest operational change was addressing the issues described in Sect. 3.5 with our custom deployment management software, described further in Sect. 4.1.

3.7 *Adaptation from Large-Scale Deployment to Sustainable Large-Scale Deployment*

After the completion of our large-scale deployment, our team was asked by MiDA to scale the deployment again, this time to over 1200 PowerWatch sensors, with the continued goal of estimating SAIDI and SAIFI and providing this data to multiple impact evaluations of the Ghana Power Compact.

Before scaling up further, a subset of our team's engineers created a company. This was not an easy decision. The academic research lab context afforded space, materials, and access to other researchers. Building a new organization has overhead, potentially taking resources away from solving deeper problems. Working across disciplines lets our team address an important data gap for economists with a new technology provided by engineers, and developing the technology as

researchers allowed us to be slow, make mistakes, iterate, and get to know the stakeholders without the pressures placed on a subcontractor delivering a product. Most importantly, freedoms enjoyed within the academy around transparency, knowledge transfer, and independence contribute greatly to our personal drive to perform this work, and the financial obligations that come with forming a company often deprioritize these goals.

We recognized, however, that this project was no longer a good fit for the goals of an academic research lab. We anticipated fewer innovations related to the sensor and deployment methodology and thus less research suitable for academic publication, while also anticipating the need to spend more time supporting the expanding deployments. Further, we recognized the need to free the project from institutional dependencies that had been too slow for the rapid pace of field work, a pace that would only increase with scale.

Since starting the company, the overhead of establishing an organization – including hiring employees, establishing accounting systems, navigating conflicts of interest, and trying to package our measurements as a product – has been significant. However, we have been more nimble and responsive when faced with challenges, have successfully scaled our deployment again, and, most importantly, have a structure in place to allow this work to exist for longer than just the length of a PhD thesis.

Pivot: Mobile Phone Sensing to Plug-in Technologies?

One original research question tested during this project was whether we could repurpose daily-use smartphones, already deployed in the hands of billions of people, as low-resolution grid reliability sensors. Leveraging the widespread ownership of cellphone devices in developing countries would allow cash-constrained utilities to improve the accuracy of reliability measurements at lower cost than widespread smart meter deployment. In fact, a primary objective of the deployment of PowerWatch, our plug-in sensor, was to provide ground-truth outage data against which the DumsorWatch app's measurements could be compared! However, as we scaled up both technologies, we saw PowerWatch succeed, and DumsorWatch underperform in two unexpected ways, one due to unanticipated participant behavior and the second due to a changing technical context. The combination of these limitations led us to prioritize PowerWatch and put DumsorWatch on the back-burner just as it was starting to return some positive results,

The first unanticipated result we observed is that DumsorWatch was uninstalled from participant phones at a higher rate than originally anticipated, even when participants were financially incentivized to keep the app installed. Most frequently, this happened automatically – for example, because the respondent reset, replaced, or lost their phone – and the respondent did not

(continued)

reinstall the app. Many participants also uninstalled their apps intentionally, for example, to save space on their phones or because they had privacy concerns from having an unknown app installed on their private device. As a result, just 3 months after the original deployment, DumsorWatch was only detecting around 10% of the outages that it was at the start. On the flipside, participants also expressed PowerWatch was less invasive than an app on the phone, potentially explaining why it regularly remained plugged in over long periods of time.

The second unanticipated action that impacted the performance of DumsorWatch was that the Android operating system changed to limit long-running background services. DumsorWatch depended on a background service that would wake the app up whenever the phone was plugged in or unplugged. This function was eliminated mid-deployment by Google in an effort to improve the user experience, since some long running apps have an outsized impact on battery life or quietly collect large amounts of data on users, burning data allocations and leaking privacy. We tried to get around this by limiting the versions of Android we recruited participants for to those that still supported background services. However, the new OS limitations ensured the sunset of DumsorWatch as an application layer technology, which led to us to reprioritize our engineering efforts to improve the PowerWatch system.

It is worth noting that our implementation of DumsorWatch successfully detected power outages! For as long as it was operational, participant phones running DumsorWatch demonstrated that uncoordinated smartphones experiencing charge state changes at the same time correlate with ground truth grid reliability measurements as provided by PowerWatch. However, based on both factors explained above, we believe the best path forward for the ideas captured in DumsorWatch is for these types of measurements to be taken as an OS level service in Android and to be aggregated as a primary crowd-sourced measurement (similar to how Google Maps captures and exposes traffic data (Google, 2009)). This pivot was made easier by the presence of PowerWatch, which could answer many of our remaining questions. While it remained a hard choice, it lets us better prioritize our team's and our funders' resources.

4 Lessons Learned

Each level of scale brought unique complexities for both engineering and operational tasks. While some of these complexities were one-time costs, many can be attributed to either the continuous nature of operating and managing a sensor deployment at scale or a context in the local culture we were not anticipating. This combination of scale and continuity stretched the administrative ability of the

university system; exploded the continuous data stream that must be maintained to manage the sensors requiring automation; amplified errors in data collection; and ultimately required us to develop automated tools to facilitate tasks that we and our field team could not handle at scale. We hope these lessons will inform future efforts to deploy continuous monitoring and evaluation systems in developing regions.

4.1 *Continuous Monitoring Requires Continuous Upkeep*

Continuous operation of a sensor network and phone application requires significant metadata and upkeep not required for 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 participants can be appropriately incentivized. All of this information needs to be effectively communicated to the field officers for debugging on an ongoing basis and updated over time because participants and their devices are in constant flux. As we describe in Sect. 3.5, maintaining a high quality of sensor deployment and implementation of our experimental design requires a systematic approach to tracking these data at scale.

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 Fig. 6.7, have two major implications for a deployment:

1. Errors introduced into the feedback loop by incorrect metadata from a survey are important and often amplified if not addressed.
2. The deployment's *state* (e.g., the location and status of each payment and device) is kept across multiple systems and is likely to become inconsistent if feedback is not automated.

For our large-scale deployment, we addressed 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 to develop custom software responsible for automatically (1) keeping *state* consistent across all databases, (2) communicating errors to the field team, and (3) implementing corrections to survey data when updates are submitted by the field team. The field team completes a set of deployment, retrieval, and debugging surveys in SurveyCTO, 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, is available to the field team through a web interface, and field team error corrections

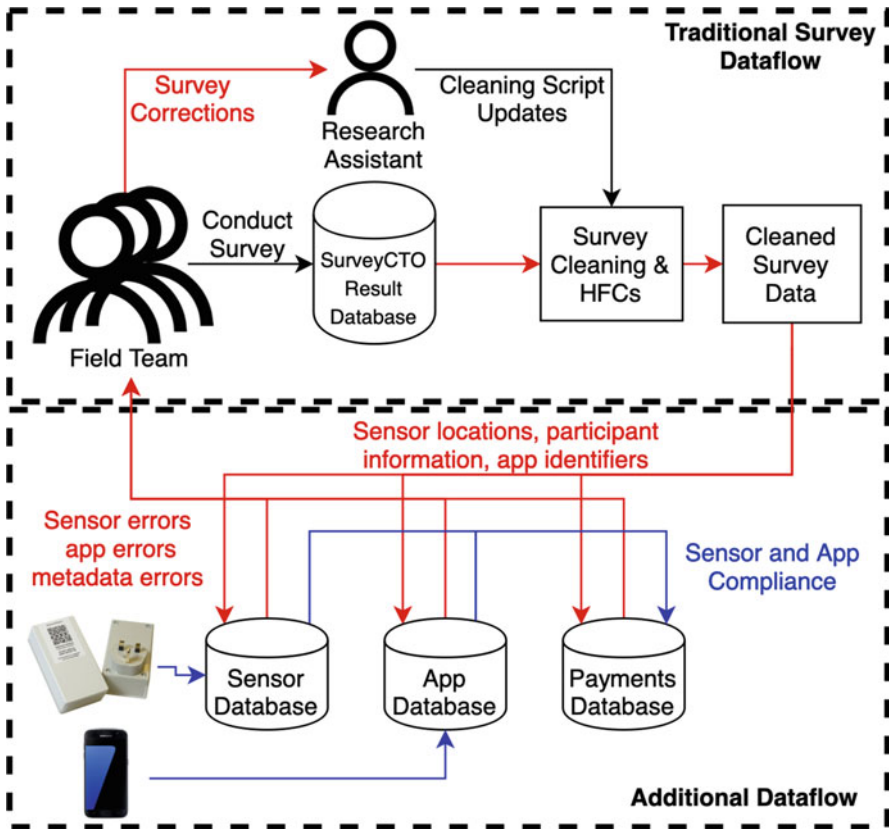


Fig. 6.7 Deployment dataflow

Notes: 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. This in turn creates more separate places where state is stored, creates more necessity to communicate this state, and amplifies the problem of errors during surveying. We implemented a deployment management system to alleviate these problems. Blue arrows represent data flows that we automated from the beginning because we anticipated their complexity before the medium-scale deployment. Red arrows represent data flows that we originally attempted to perform manually and that we now automate or facilitate with a deployment management tool

are communicated back to the software through a final correction survey. The deployment management software is represented by the red arrows in Fig. 6.7.

This architecture makes the deployment significantly easier to manage. 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.

4.2 Global Solutions May Miss Local Context

Several times in our deployment we were forced to consider tradeoffs between using technology and services developed and operated locally and 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 airtime incentives to participants. Unsurprisingly, we found local service providers were more likely to provide high-quality service in Ghana compared to US-based companies, which had 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 was prepared to bill and support our enterprise accounts. Therefore, to continue to scale, we went back to evaluating mobile virtual network operators (MVNO) and global airtime top-up providers, aiming to find companies with demonstrated experience in Ghana and similar geographies. After evaluating several MVNOs and airtime top-up firms, we found a set of global companies that provide a good mix of technical maturity, experience in handling enterprise customers, and reliable service in Ghana and other countries.

4.3 University Lacks Financial Agility

One of our primary organizational problems was the inability to pay for the various services necessary to perform our deployment. This was not for the lack of available funding but a mismatch with the administrative capacity in academia.

While our university policy dictates a single-day turnaround on wire transfers, in practice this often took over 15 days. Setting up subcontracts with companies, especially companies with which the university had never contracted before, often took months, and our deployment required numerous contracts because our technology relies on external service providers. As a result, 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 a deployment, members of the research team still frequently resorted to using private financial resources, maxing out 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 personal credit card.

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




Fig. 6.8 Participant perception of sensors (a) Was DumsorWatch or PowerWatch preferred? (b) Are you likely to download DumsorWatch when . . . ?

Ghana early in the deployment so they could better handle later delays in payment from the university. Even still, it would be difficult to imagine the deployment running at its described pace without the research team extending personal credit.

4.4 Technology Usage Patterns Impact Design

Our system depends on participants to download apps and install sensors 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 large-scale deployment. The results of this survey proved surprising and critical for the design of the next level of scale.

We asked participants what they thought of the sensors. Figure 6.8a shows that participants liked both PowerWatch and DumsorWatch, with a slight preference for PowerWatch, 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 scarce and highly valued.

We then explored how incentives influence participation. Figure 6.8b shows that real-time information about local outages was valued highest, indicating a strong

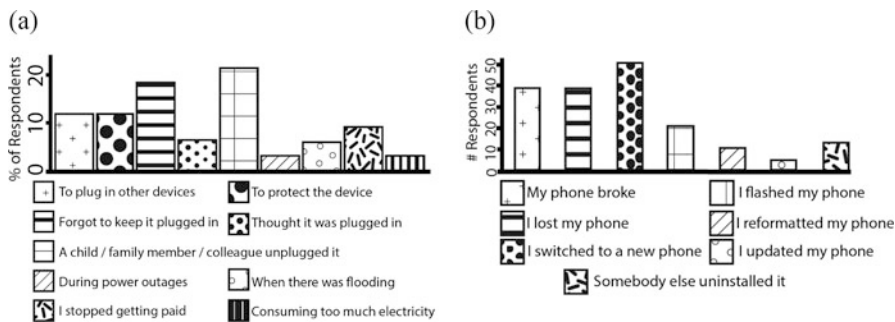


Fig. 6.9 Participant-reported reasons why sensors were uninstalled (a) PowerWatch (b) Dumsor-Watch

local desire for energy-reliability data and suggesting 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 6.9a 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 Fig. 6.9b, 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 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.

The results of this survey were important for improving system performance and would have been even more effective if we had run it earlier. The successful deployment of any technology hinges on its suitability to the local context. Developing an early understanding of the potential risks, for example, through qualitative surveys or focus groups, can help inform the design of a technology while there is still ample scope to do so.

5 Conclusion

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 not due to 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 these hurdles, we conducted a largely successful deployment that met 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.⁵ In reaction to specific pain points at larger scales, we developed meta-tools, not to replace the human links but to assist them.

We hope that by identifying and describing these meta-tools, as well as our broader collection of lessons learned, we lower the barrier to entry for conducting similar-scale deployments. This is a goal worth working towards. Insights gathered from direct deployments of sensors around the world will be critical for achieving deeper understanding how critical life sustaining systems operate and, eventually, how they can be improved. To realize this, the development engineering community will have to work to reduce the costs and complexities of performing deployments. Steps towards this may include continually sharing lessons learned, making software meta-tools available and open-source and setting up incentive systems within academia to help engineering researchers value innovation in methodology, reliability, and scale-ability as highly as innovation in any other more traditional high-order subsystems.

Discussion Questions

1. When planning an interdisciplinary project, how should you deal with differences in timelines, goals, or expectations between stakeholders from different disciplines or with different incentives?
2. When planning a real-world deployment, how do you balance and prioritize your expected contributions to science (direct professional benefit) and your expected contributions to the population under study (direct societal benefit)?
3. What do you do when a project unexpectedly requires a skill you have not received training in?

⁵ This work was supported in part by the Millennium Challenge Corporation through 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. We additionally thank CEO Martin Eson-Benjamin, CFO John Boakye, Director Dr. Kofi Marfo, Sharon Abena Parku, and the entire team at the Millennium Development Authority (MiDA); Jeff Garnett, Hana Scheetz Freymiller, and the rest of our partners at the Millennium Challenge Corporation (MCC); Karen Notsund; Nana Bonku Eshun; Kobina Aidoo; Ivan Bobashev; Aldo Suseno; Revati Thatte; Lane Powell; and staff from the Electricity Company of Ghana, SMEC Ghana, and ESB International. Finally, we thank the Field Supervisors and Field Officers in Accra who put in many hours of hard work to implement, improve, and maintain our deployment.

4. How do you assess the efficacy of a service provider when you will be the first to empirically test their capacity to scale in a given market?
5. After your research study is complete, what stakeholders will manage your deployment? How will you ensure the infrastructure is useful and usable for them? What is your handoff plan, both for short-term management and long-term maintenance?
6. Where is information stored in your system, what are the sources of this information, and what other systems depend on this information? Are systems in place to ensure that information is correct and consistent for different data consumers?

References

- Acakpovi, A., Abubakar, R., Asabere, N. Y., & Majeed, I. B. (2019). Barriers and Prospects of Smart Grid Adoption in Ghana. *Procedia Manufacturing*. In *2nd International Conference on Sustainable Materials Processing and Manufacturing*.
- Ackah, C. (2015). Electricity insecurity and its impact on micro and small businesses in Ghana. In *Study presented at Institute of Statistical, social and economic research (ISSER) workshop*. University of Ghana, Legon.
- Aidoo, K., & Briggs, R. C. (2019). Underpowered: Rolling blackouts in Africa disproportionately hurt the poor. *African Studies Review*, 62, 3.
- Banuenumah, W., Sekyere, F., & Dotche, K. A. (2017). *Field survey of smart metering implementation using a simple random method: A case study of new Juaben municipality in Ghana*. IEEE PES PowerAfrica.
- Bowers, B. (1988). Social benefits of electricity. *IEE Proceedings A (Physical Science, Measurement and Instrumentation, Management and Education, Reviews)*, 135, 5.
- Clerici, C., Taylor, M. S., & Taylor, K. (2016). *Dumsor: The electricity outages leaving Ghana in the dark*. Al Jazeera.
- Electricity Company of Ghana. (2019). *Prepayment Metering*.
- Energy Commission Ghana. (2019). *National Energy Statistics 2000–2019*.
- Gertler, P. J., Lee, K., & Mobarak, A. M. (2017). *Electricity reliability and economic development in cities: A microeconomic perspective*.
- Google. (2009). *The bright side of sitting in traffic: Crowdsourcing road congestion data*. Official Google Blog.
- Google. (2018). Google Play services discontinuing updates for API levels 14 and 15. (Dec 2018). <https://android-developers.googleblog.com/2018/12/google-play-servicesdiscontinuing.html>
- Hamidu, M., & Sarbah, A. (2016). An assessment of the electrical energy needs of beauty saloon industry in Ghana. *American Journal of Management Science and Engineering*, 1(2).
- He, J., Cheng, L., & Sun, Y. (2007). Transformer real-time reliability model based on operating conditions. *Journal of Zhejiang University-Science A*, 8(3).
- IEEE Guide for Electric Power Distribution Reliability Indices. (2001). IEEE Std 1366, 2001 Edition (2001), i–. <https://doi.org/10.1109/IEEESTD.2001.94438>
- Klugman, N., Rosa, J., Pannuto, P., Podolsky, M., Huang, W., & Dutta, P. (2014). Grid watch: Mapping blackouts with smart phones. *ACM HotMobile*, 14.
- Klugman, N., Correa, S., Pannuto, P., Podolsky, M., Taneja, J., & Dutta, P. (2019). The open incentive kit (OINK): standardizing the generation, comparison, and deployment of incentive systems. *ACM ICTD '17*.

- McCollum, D., Gomez Echeverri, L., Riahi, K., & Parkinson, S. (2017). *SDG7: Ensure access to affordable, reliable, sustainable and modern energy for all* (pp. 127–173). International Council for Science. <https://doi.org/10.24948/2017.01>
- Millennium Challenge Corporation. (2008). *Guidance on Common Indicators*.
- Millennium Development Authority. (2014). *Power distribution feasibility studies, Ghana Phase II Feasibility Assessment*.
- Millennium Development Authority. (2018). *Ghana monitoring and evaluation plan*.
- Nunoo, S., & Ofei, A. K. (2010). *Distribution automation (DA) using supervisory control and data acquisition (SCADA) with advanced metering infrastructure (AMI)*. *IEEE Conference on Innovative Technologies for an Efficient and Reliable Electricity Supply*.
- Quayson-Dadzie, J. (2012). *Customer perception and acceptability on the use of prepaid metering system in Accra west region of electricity company of Ghana*. Ph.D. Dissertation.
- SurveyCTO. (2019). Because your data is worth it. <https://www.surveyccto.com/>. Accessed on 03/15/2019.
- Sustainable Energy for All. (2018). *Sustainable energy for all: Monitoring, evaluation, and learning framework*.
- Taneja, J. (2017). *Measuring electricity reliability in Kenya*. Technical report. Working paper. Available at <http://blogs.umass.edu/jtaneja/files/2017/05/outages.pdf>
- The GridWise Alliance. (2013). *Improving Grid Reliability and Resilience: Workshop Summary and Key Recommendations*. <https://www.energy.gov/sites/prod/files/2015/03/f20/>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

