

# Fine-Grained Remote Monitoring, Control and Pre-Paid Electrical Service in Rural Microgrids

Maxim Buevich <sup>†</sup>      Dan Schnitzer <sup>‡</sup>      Tristan Escalada <sup>\*</sup>

Arthur Jacquiau-Chamski <sup>\*</sup>      Anthony Rowe <sup>†</sup>

<sup>†</sup> Electrical and Computer Engineering Department, Carnegie Mellon University

<sup>‡</sup> Engineering and Public Policy, Carnegie Mellon University

<sup>\*</sup> EarthSpark International

{mbuevich,dschnitz,agr}@andrew.cmu.edu

{tristan,arthur}@earthsparkinternational.org

**Abstract**—In this paper, we present the architecture, design and experiences from a wirelessly managed microgrid deployment in rural Les Anglais, Haiti. The system consists of a three-tiered architecture with a cloud-based monitoring and control service, a local embedded gateway infrastructure and a mesh network of wireless smart meters deployed at 52 buildings. Each smart meter device has an 802.15.4 radio that enables remote monitoring and control of electrical service. The meters communicate over a scalable multi-hop TDMA network back to a central gateway that manages load within the system. The gateway also provides an 802.11 interface for an on-site operator and a cellular modem connection to a cloud-backend that manages and stores billing and usage data. The cloud backend allows occupants in each home to pre-pay for electricity at a particular peak power limit using a text messaging service. The system activates each meter within seconds and locally enforces power limits with provisioning for theft detection. We believe that this fine-grained micro-payment model can enable sustainable power in otherwise unfeasible areas.

This paper provides a chronology of our deployment and installation strategy that involved GPS-based site mapping along with various network conditioning actions required as the network evolved. Finally, we summarize key lessons learned and hypothesis about additional hardware that could be used to ease the tracing of faults like short circuits and downed lines within microgrids.

**Keywords**—Microgrid, Deployment Experience, Electricity Distribution

## I. INTRODUCTION

Nearly 1.3 billion people, mostly living in less developed countries, do not have access to electricity [1]. Thanks to advances in technology, decreases in cost, and improvements to telecommunications and supporting infrastructures, microgrids are becoming an increasingly common means for provisioning electricity in remote areas [2].

In Haiti, less than 25% of the population has access to electricity [3], leaving the vast majority of the country in energy poverty. Although more than 30 microgrids have been developed in Haiti since the mid-1980s, few of them continue to operate sustainably. Powered completely by diesel generators and operated by local municipalities, some of these microgrids operate for as few as 150 hours per year. The unsustainability of these projects and their failure to reliably deliver energy services is symptomatic of the reasons for many microgrids to fail: (1) poor cost recovery due to underpriced

tariffs and weak revenue collection systems (2) inability to limit customer usage to levels below generation output capacity (3) ability to meter usage at a fine enough granularity to quickly diagnose problems and detect theft. In this paper, we present the design and a case-study of a system that addresses these challenges.

Many microgrid developers do not utilize meters to track household electricity consumption and issue corresponding bills because meters add cost to systems that are designed to serve very low-income households [4]. Instead, these microgrid operators charge customers a fixed monthly fee. Unfortunately, fixed monthly payments are often priced at levels far below the cost of the energy used by the customer during that time. Operators also face challenges collecting monthly payments on time, as impoverished populations often do not have steady income streams to pay ongoing expenses. Pre-payment resolves the issue of cost recovery by addressing the issues of both poor pricing and poor collection. Customers are empowered to align their preferences with their budget constraints by purchasing electricity when funds are available and when they need it most.

Imposing load limits on customers is crucial for microgrids where consumption may reach the maximum output of the generation source [5]. Under such conditions, the microgrid is forced into a brownout, precluding all users from access to electricity. This situation can arise from a number of circumstances. On some microgrids where the system is sized to provide lighting with CFL bulbs, customers using incandescent light bulbs can drive the microgrid into a brownout. In Bhutan, there is a well-documented case of microgrids browning out when large numbers of customers concurrently power their electric rice cookers in the evening [6]. In response to the threat of brownouts, microgrid operators can employ a number of load management techniques (exhaustively described in [4]), from simple bans on certain loads, to circuit breakers on customer households, to sophisticated meters with programmable relays. We propose using local monitoring on each of the meters along with a global view of the network with low-latency metering and control to address these concerns.

Theft is another issue to contend with on remote microgrids [5]. Most often, theft is carried out by making an unauthorized connection to the microgrid distribution line. In other cases, theft is carried out by authorized customers who bypass their

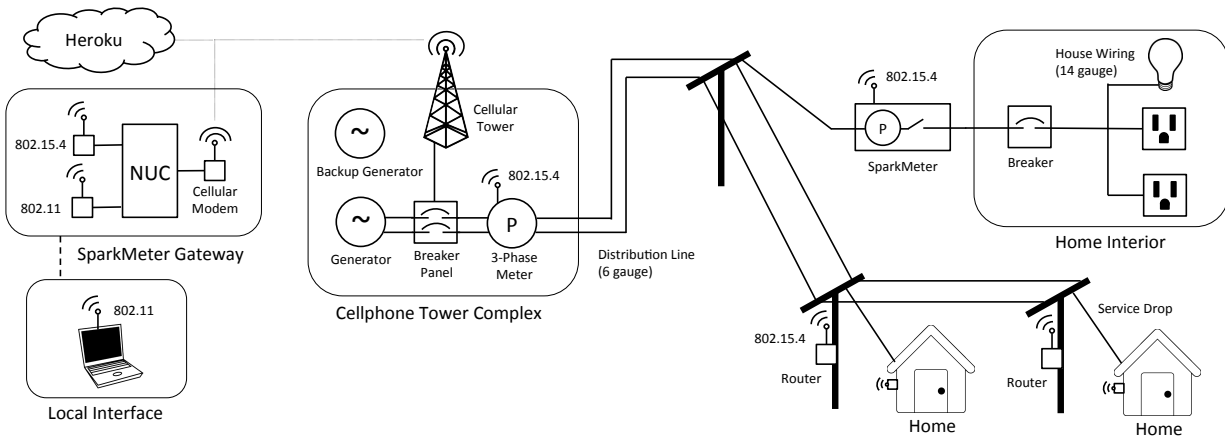


Fig. 1. System Architecture



Fig. 2. Haiti Satellite View




Fig. 3. 3D Model of Deployment

meters. Monitoring theft is therefore difficult, especially on systems that serve a few hundred households. Theft is most often dealt with through strong local institutions that can impose a credible threat of penalty. However, penalties are often unenforced, and theft persists on many systems. Our system uses time synchronized sampling of power data across the network to aid in loss estimation.

In 2012, EarthSpark International developed a microgrid in the town of Les Anglais, Haiti to demonstrate a sustainable model for microgrid operations encompassing pre-payment, load control and theft detection. Beginning operations on November 20th, 2012, the microgrid provided power to 14 households 24 hours per day. It shares power with the local cellular base transceiver station from a 12 kVA diesel generator. The cellular tower typically uses between 2 and 5 kVA, leaving ample headroom on the diesel generator to provide power for basic loads such as lighting, cell phone charging and entertainment to a number of households in the community. The most basic level of service provided in our system has a 30W per house limit while the highest consumption level is a 360W limit that can be used for refrigeration. Prior to the implementation of the wireless metering system, the EarthSpark microgrid relied on manual meter reading to update pre-paid account balances on a daily basis, miniature circuit breakers to limit loads, and a strong local presence to deter theft. While operations have been smooth since the inauguration of the grid, the manual nature of the system is not well-suited to

scaling up to the remaining 400 households in Les Anglais to be connected to the microgrid. A system that enables pre-payment on such a scale would benefit from: automatic energy readings to update account balances, automatic actuation in response to load limit violations and account balance status, and remote monitoring to enable theft detection.

In this paper, we discuss the design, architecture and initial deployment experiences gained by updating and expanding the 14-customer manually controlled microgrid with wireless smart meters on 38 new customers. This effort included the design of custom low-cost wireless energy meters, the development of a meter networking protocol and involved instrumentation for efficient deployment and debugging. The metering hardware communicates back to a local gateway using IEEE 802.15.4 radios and then to a cloud billing management system using an EDGE cellular connection. Wireless meter communication is achieved through a TDMA networking protocol that combines elements of the Flash Flood protocol [7] and the Low-Power Wireless Bus protocol [8]. We chronicle (and critique) our deployment strategy which involved GPS-based packet reception mapping of the town as part of a pre-planning step for tiling forwarders. This also involved a significant tuning effort one month after the system was installed. We then provide an evaluation of the system's performance in terms of packet yield, message delay (critical for load management) and data fidelity. Finally, we discuss lessons learned from the deployment and introduce a hardware extension concept




Fig. 4. Hardware Photos

that we believe will drastically simplify future installations by including automatic distributed fault tracing.

## II. RELATED WORK

In this section we examine existing research related to rural microgrid distribution systems. We first discuss related projects, followed by relevant support technologies.

### A. Rural Microgrid Projects

Microgrids can be effective means to provide access to electricity in areas of the world where it is financially, physically or institutionally difficult to extend the reach of the central grid [5]. They are developed by government agencies, private developers and NGOs around the world. Microgrid developer portfolios range from small deployments such as DESI Power’s four village microgrids in Bihar, India powered by biomass gasification [9] to the West Bengal Renewable Energy Development Agency’s 18 microgrids powered primarily by solar PV in the Sundarbans [10] to Nepal, which powers 59,000 households with over 300 micro-hydro systems [11].

Microgrids vary widely in the services they provide. Some microgrids are designed to provide only the most basic energy services - such as lighting in the evening - while others are designed to power cold storage and other productive commercial or agricultural loads such as mills and irrigation pumps [12].

Unfortunately, many microgrids fall into non-functioning states due to any one of a number of factors, including low levels of tariff collection, poor maintenance, customer over-usage (which causes brownouts), and unmet growth in demand [5]. In recent years, attempts at modeling microgrid operations have shown that a number of interventions can improve microgrid sustainability, such as the use of renewable energy to improve cost-effectiveness [13] [14] [15], energy efficiency [15], and the use of demand side management (DSM) strategies and technologies [4] [6].

In recent years, microgrid developers are turning to advanced pre-paid meters to solve problems of customer over-use and poor tariff collection. Companies such as Devery [16], Circutor [17], Inensus [18], Powerhive [19] and Gram Power [20] have each developed such metering systems targeted at microgrids serving low-income customers. While they are in their nascent stages of pilot projects or field trials, these

systems offer significant advantages over both conventional metering systems and less sophisticated pre-paid meters such as those made by Conlog [21]. Those advantages include greater flexibility in billing, remote system monitoring, price-responsive DSM, load-shedding, and dynamic power- and energy-limiting [4].

### B. Wireless Energy Management

Historically, it is not uncommon for grid-level energy management systems to use wireless communication for metering. Many of the new generation of smart meters in use today have PLC or wireless backhauls. Some even have 802.15.4 capabilities for communication with appliances in the home. These systems are typically expensive and not optimized for the challenges found in microgrids because they are built to 100A (or greater) standards and are designed for use over robust distribution or communication systems. Furthermore, while forward-looking demand response programs are investigating many of these topics, none have deployed systems with fine-grained control over a non-trivial set of homes.

Multiple research groups and companies have developed similar wireless energy monitoring solutions for buildings. The MIT Plug [22] provides users with power and sensor information by means of a smart surge protector. In [23], the author’s present experiences using the ACme wireless plug sensor in an office environment. This hardware platform is quite similar in nature to our solution except that we provide the surrounding eco-system for micro-grid management. Companies like Tendril Inc. [24], AlertMe [25], Trilliant [26] and GreenWave Reality [27] provide solutions that allow home owners to install and monitor their own energy usage. In some cases [24], they are even working with utility companies as part of demand response programs. These are similar technologies, but the installation and use-cases are quite different since micro-grids often require low-latency active management of load.

## III. SYSTEM ARCHITECTURE

In this section we provide a brief overview of the design of our electricity metering and control architecture shown in Figure 1. The main components consist of a cellular phone tower complex which houses two diesel generators that feed the electrical distribution network. Each home is connected to the distribution network using an 802.15.4 wireless metering



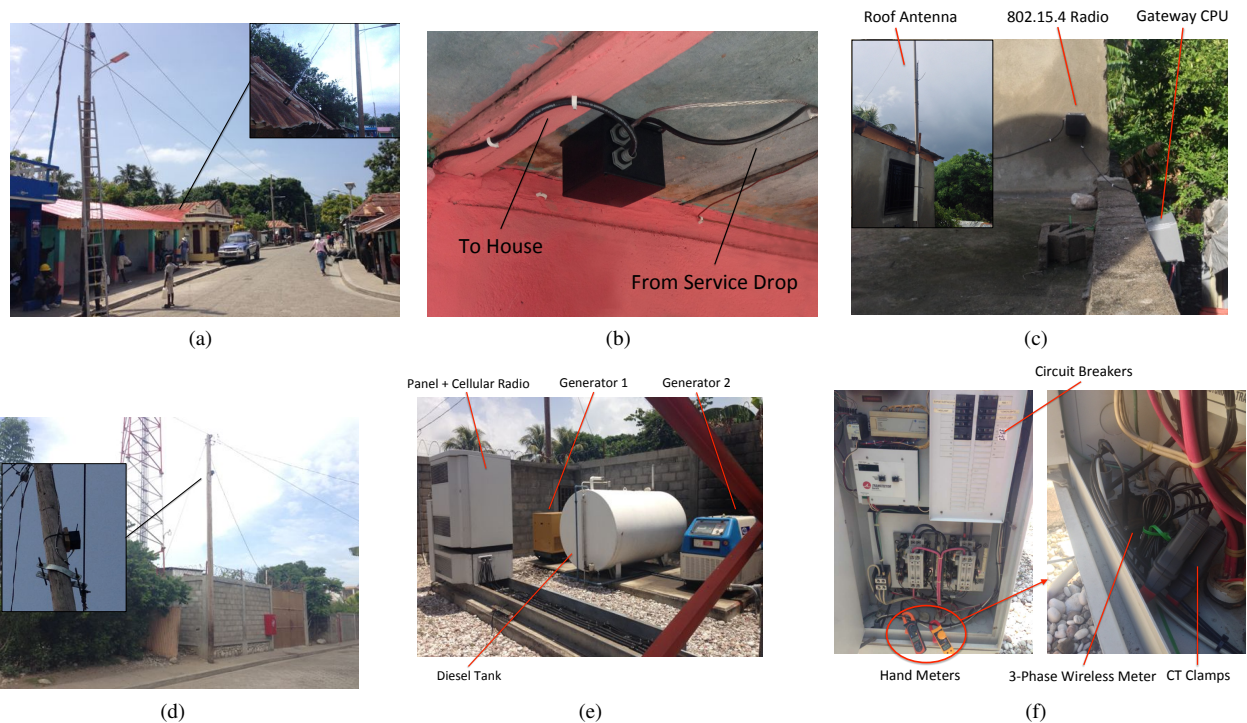


Fig. 5. Deployment photos: (a) street with forwarder and highlighted service drop (b) sparkmeter on outside of home (c) gateway hardware (d) cellular tower compound with forwarder (e) cellular tower interior and (f) main electrical panel

and control device called a SparkMeter. The SparkMeters communicate over multiple hops through powered forwarders mounted on select electrical poles back to a gateway housed in the local management shop. The gateway is responsible for on-site data logging and low-latency management tasks. The gateway also acts as a local WiFi hotspot to serve a management web interface that is used by operators to monitor meters, balance billing and configure individual devices. The gateway has a cellular modem that uses the local EDGE network to communicate with a cloud management backend that stores aggregate data and has a richer set of configuration options for managing multiple connected grids.

Figure 2 shows a satellite view of Haiti, as well as other island nations of the Greater Antillean archipelago (Cayman Islands and Puerto Rico not shown). A 3D model of the physical topology of our deployment in Les Anglais, Haiti can be seen in Figure 3. The cellular tower is shown to the far right. The building shades (colors) indicate different customer service levels.

#### A. Hardware Components

The field metering hardware is built upon the FireFly sensor networking platform and consists of a 3-phase meter, the SparkMeter nodes and the forwarder units. Each component uses an ATmega128RFA1 micro-controller with 16KB of RAM, 128KB of flash memory, running at 16MHz with an integrated on-chip 802.15.4 radio for communication.

1) *SparkMeter*: The SparkMeter, shown in Figure 4(a), is the interface between the distribution system and each home's internal wiring. Figure 5(b) shows the board deployed inside a waterproof NEMA enclosure. Current is measured using

an ACS712 Hall effect sensor and voltage is read directly using a voltage divider. The output of the Hall effect sensor is passed into two separate amplification circuits, one for high-gain (loads less than 100W) and one for low-gain (loads less than 1650W). Each input is fed into different channels of the ATmega's onboard 10-bit ADC and sampled at 2KHz to compute frequency, RMS voltage, RMS current, true power, apparent power, power factor and energy. The multi-gain inputs allow for high-resolution monitoring of the small loads that are common to microgrid environments (some customers are limited to 30W total per home). The SparkMeter can switch the load on and off using a 15A 240VAC mechanical relay which is more than sufficient for an entire home in rural areas. The power supply for the board produces 5V for the sensors, 3.3V for the micro-controller and 48V for the relay that is switched using a FET. To reduce node costs, we opted to use a PCB antenna and no power amplifier (in retrospect this was a poor choice).

The firmware running on the SparkMeter is responsible for four main tasks. First, it must sample the raw voltage and current values at 2KHz and integrate voltage and current to produce true power updates. Second, it runs the communication protocol stack described in Section IV. It also periodically stores status and error logging runtime statistics to EEPROM. Finally, it runs a power-limit enforcement task that checks once per second to make sure the meter does not exceed an admin-defined load profile. If the power consumption exceeds these limits, the node will disable power and every 30 seconds switch on the relay to check if the limit is still being violated. As part of the load profiles, we change the monitoring window based on the allowed power level. This programmability prevents devices with large transients from constantly blowing breakers



like in many other limiting systems. The MAC address of each node is used to delay when the meters enables power after an outage to avoid large inrush currents that can cause additional brownouts.

2) *Forwarder Nodes*: The forwarder nodes consist of an ATmega128RFA1 with the addition of a CC2591 amplifier that allows for longer-range communication across the network. The main forwarder board shown in Figure 4(d) has an external SMA connector and can be powered from a standard USB power adapter. These nodes are packaged in waterproof NEMA boxes and installed at the tops of utility poles, with direct access to main voltage lines and elevated line-of-sight communication with neighboring forwarders. With the addition of the CC2591 amplifier, these nodes have 3-4 times the range of the SparkMeter nodes and are able to effectively provide coverage for wide areas.

3) *3-Phase Meter*: The 3-Phase meter board was designed around the ADE7878 3-phase power metering IC. The ADE7878 is highly accurate and used in many industrial metering applications. It connects as a sensor interface to one of the power amplified ATmega128RFA1 nodes. This board requires at least one voltage input and can support up to three different current transformer inputs. Figure 5(f) shows the 3-phase meter installed in the breaker box at the cellular tower and connected to the two input phases of the distribution system. Metering the generator directly allows the system to monitor total microgrid load (even if some SparkMeters are down) and enables the system to perform real-time comparisons of output versus metered load for the purpose of theft detection.

4) *Gateway*: The gateway is located inside our local partner’s shop and consists of an Intel NUC mini-PC connected to a MOXA cellular modem and a router node shown in Figure 5(c). In our initial deployment, the NUC contains a 1.8GHz Intel Core i3 processor with 16GB of RAM and a 64GB solid-state drive to simplify remote monitoring and patching. However, our software is also able to run on more resource constrained embedded platforms like the Raspberry Pi which would significantly lower costs for future deployments. The gateway has a 20ft coax extension cable which runs to a roof-mounted antenna for improved connectivity with forwarder nodes.

## B. Software Architecture

The SparkMeter system software architecture is divided between the local gateway architecture and the cloud-based architecture as shown in Figure 6. Locally, software runs on an Intel NUC mini-PC connected to a 802.15.4 USB node and a GSM/GPRS MOXA modem, referred to as the gateway. An SSH tunnel is used to connect over EDGE to a Heroku instance.

1) *Local Gateway Software*: The local software serves the following core functions: (1) Manages the sensor network, (2) Stores real-time packet data, (3) Tracks customer account balances, (4) Synchronizes data with the cloud software, and (5) Provides a local interface for microgrid operations. Local microgrid operations include pre-payment of customer electricity credits, and monitoring of real-time data.

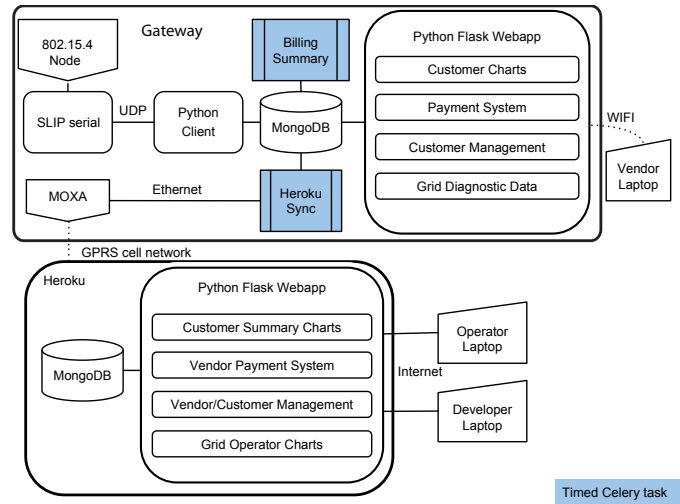


Fig. 6. Software Architecture

**Server Task:** The gateway is where most of the main processing on the microgrid is hosted. Packets are received by the gateway and transferred to the gateway through UDP/SLIP. Packets are processed by a Python client and written to a MongoDB database.

All data processing is controlled using Celery, an asynchronous task queue, with a RabbitMQ server as its backend. The Celery tasks operate on data stored in MongoDB. Once written to the database, data is handled by the tasks described below.

**Billing Task:** Every minute, a billing task is run that summarizes all of the readings received since the last task-run. This summary reading includes time-stamped values for power consumed, power factor, frequency, voltage, and current. It also includes the amount of energy consumed in the last minute and the cost of that energy. The one-minute energy costs are then deducted from the customer balances. If a customer’s balance becomes negative, then a command is sent to that customer’s meter to be disabled.

**Cloud Synchronization:** Every fifteen minutes, a synchronization task is run that uploads all of the saved summary readings to Heroku via the GSM / GPRS MOXA modem. This task also checks for any updates in Heroku to see if a grid operator has added any credits to vendor accounts, changed global settings, or made other operational changes from the remote interface.

**Local Interface:** A Python web application serves a website that acts as the local interface. The pre-payment electricity credit vendors can log into the website via WiFi or a direct Ethernet connection to the gateway from their laptops, tablets or smartphones. Once logged in, the vendor is able to see all customer accounts and their current status and usage.

When a customer wants to add funds to their balance, they make a payment (cash or mobile money) to the vendor. The vendor enters the payment information through the interface that will increase the customer’s balance, and immediately turn back on their electricity if they had run out of credits. The local interface also provides a webpage of charts displaying the health of the grid in real time. If an issue arises with the

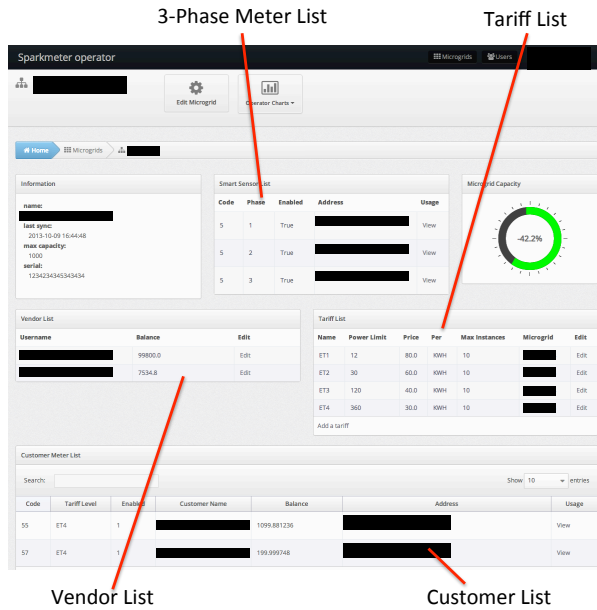


Fig. 7. Remote Interface

grid that cannot be diagnosed with the one-minute summary data displayed in the remote interface, then the vendors can login locally to help the operator.

2) *Heroku Remote Software*: The remote, cloud-based software serves the following core functions: (1) Stores summary data, and (2) Provides a remote interface for microgrid operations. Remote microgrid operations include pre-payment of vendor electricity credits, defining tariff levels, monitoring system sensors such as the 3-phase meter, and monitoring and controlling customer meters.

The Heroku instance runs on the same codebase as the gateway, but when run in the cloud, a different set of functionality is enabled. The main interface dashboard is shown in Figure 7. An operator has access to the key functions listed below.

**Vendor Management**: Local vendors are responsible for selling electricity credits to microgrid customers on a pre-paid basis. The remote interface allows the operator to add or remove vendors, and to add electricity credits to their accounts upon receipt of payment.

**Tariff Level Management**: The remote interface enables microgrid operators to arbitrarily define an unlimited number of tariff levels. Each tariff level has a corresponding power limit and price and may be assigned to any customer. The interface displays the maximum number of instances for each tariff level given the total capacity of the microgrid, which keeps operators from "over-subscribing" and putting the microgrid at risk of aggregate over-consumption.

**System Sensor Monitoring**: The SparkMeter system enables microgrid operators to deploy any number of system sensors, such as multi-phase meters on the generation system feeder, meters on the overhead distribution system, and other sensors on inverters or charge controllers if present. The system deployed in Les Anglais uses a 3-phase meter monitoring the two main feeder lines powering the distribution system.

**Customer Monitoring and Control**: Microgrid operators have access to lists of all customers on each microgrid system. Selecting an individual customer, operators can view summary charts for status variables such as frequency and voltage, consumption charts for power and energy, and financial charts for customer account balances, hourly cost of energy used, and bill payments. From each customer webpage, operators can assign customers to their appropriate tariff level. If so desired, operators can also actuate individual customer meters.

## IV. NETWORK DESIGN

Microgrids can easily have tens to hundreds of controllable end-points each of which must be able to report and respond to commands within seconds to avoid brownouts. It is difficult to quantify the exact timing requirements for load balancing and brownout prevention in our target diesel generator, however in practice we see our generator able to operate above nominal capacity for about 10-20 seconds before it browns out or engages safety overrides.

### A. Research Challenges

In order to automate and connect a large number of meters, we face the following networking challenges:

- 1) **Node Placement**: The provisioning of the network requires careful placement of gateway and forwarder nodes. Peripheral nodes are unreliable when coverage is poor while overprovisioning leads to increased network traffic and contention.
- 2) **Constrained Uplink**: Cellular and wired backhauled are both expensive and often intermittent.
- 3) **Reliability**: Maintaining meter connections is critical for load management and billing accuracy.
- 4) **Visibility**: In order to diagnose and dispatch maintenance work the ability to monitor and evaluate routes within regions that have poor connectivity is critical.
- 5) **Security**: The system must be designed to robustly handle issues of tampering and fraud.
- 6) **Scalability**: Our target deployment sites often have hundreds to thousands of homes. The communication protocol must elegantly scale in a predictable manner to accommodate many nodes.
- 7) **Latency**: The decision to disable a power-meter must be dispatched from the cloud service (not locally at the meter) in case users replenish their charges before the latest quota expires. Maintaining disruption-free service is paramount to keeping users engaged. If a node does not report or suffers a long deactivation delay, the user is able to consume power for free. Long activation delays also prevent usage and negatively impact the effectiveness of the pre-payment model. Responsive networking on the order of tens of seconds has a direct impact on the system's overall quality and operation costs. Since reading data from each meter can take on the order of 10s of seconds, it is important that sampling of power data is synchronous. Synchronous sampling is critical for supporting theft detection since loads may change across a network collection cycle.

Given these design requirements, we decided to base our networking protocol, called mPCF, on a TDMA solution similar to the Low-Power wireless bus protocol described

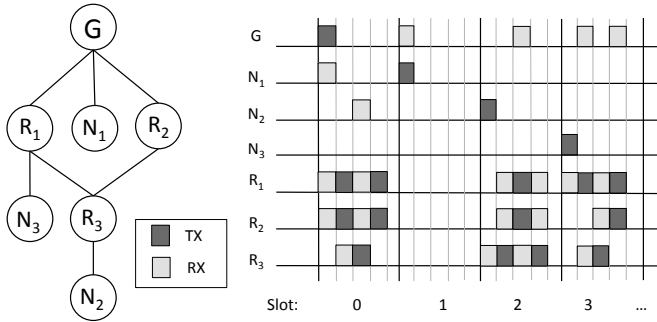


Fig. 8. Protocol Transaction Timeline

in [8]. In [8] the authors use the Glossy [28] time synchronization primitive to establish the beginning of TDMA communication slots. Each node in the system is then given a scheduled collision-free time to reply with a message that again uses Glossy to return to the gateway. Unfortunately, Glossy requires that all messages be identical so that receivers capture time-synchronized transmissions from two different nodes as a single message can still be decoded. Glossy is powerful in terms of performance, but this restriction means that forwarders cannot append unique route or diagnostic data at each hop to trace the exact trajectory through the network. To avoid this problem, we leverage the Flash Flood protocol [7] that capitalized on the radio capture effect to flood messages across the network. Using Flash Flood instead of Glossy for flooding, each end-node responds to a downstream start-of-frame message from the gateway using a fixed TDMA schedule. Replies are then flooded back again using Flash Flood within a fixed window time. By using fixed hop-counts and low-level time stamping (described in FTSP [29]), each meter is able to estimate and subtract message propagation delays so that it can accurately transmit at the beginning of its time-slot, even if it received the sync message multiple hops into the network.

Our network topology consists of many low-powered meter nodes services by a small number of high-powered forwarder nodes. Flash Flood is used to provide meter nodes with contention-free one-hop communication to the nearest forwarder. This combination of techniques allows our system to benefit from many of the timing and performance advantages of low-power wireless bus while at the same time allowing us to collect real-time route and packet-loss diagnostics from Flash Flood. Flash Flood strives to balance load with contention such that the capture effect still allows concurrent radio transmissions. We show in Section IV-D that when the density is relatively low, this approach provides good packet reception rates while being extremely robust in terms of routing.

### B. mPCF Protocol

Figure 8 outlines an example transaction timeline of the mPCF protocol. The gateway downstream message contains a header used for time synchronization, a nonce for seeding packet encryption, TDMA schedule tuning and any pending control messages for the meters. In Figure 8, transmitted messages are shown next to each node on the timeline as a dark grey box while received messages are light grey. The




Fig. 9. Impact of forwarder density on PRR

timeline illustrates the first three slots of a communication round assuming the connectivity topology graph, where edges represent viable communication links. Nodes  $N_1$  through  $N_3$  are end-point meter nodes and  $R_1$  through  $R_3$  represent forwarder nodes. In slot 0, the gateway transmits its beacon which is received by  $N_1$ ,  $R_1$  and  $R_2$ .  $N_1$  receives the message uses the  $TTL$  and  $MAX-TTL$  values in the header to compute its offset from the start of the TDMA cycle which in this case is 0.  $N_2$  on the other hand does not receive the master beacon until the third mini-slot within slot 0 and hence must subtract that delta from its next wakeup. This is accomplished by subtracting the packet transmission time  $T_p$  times the number of hops based on  $TTL$  from the slot length  $T_s$ .

### C. Security

Each packet is encrypted and authenticated using the ATmega's built-in 128-bit AES encryption hardware. There is a single shared private key that is set for each network at deployment time. We use the cipher block chaining message authentication code (CBC-MAC) capability provided by the hardware for message integrity. Since data is frequently reported by nodes, if they do not respond over a long enough interval, a service request is generated for them to be manually inspected to track physical tampering. In the future, we also intend to track the 3-phase meter output along with the sum of all SparkMeters in order to detect theft.

### D. mPCF Parameters

The mPCF header transmitted each cycle by the gateway contains information about how many slots are in each TDMA frame, the length of each slot, the maximum number of hops, the current hop count, and a time-of-day timestamp that can be used for local calendar scheduling. Each node that receives a beacon updates its internal state based on the latest TDMA definitions, making the protocol highly elastic. If the network grows in terms of nodes, more slots can be added per cycle from the gateway at runtime. Likewise, if the network grows in terms of hop-count, the gateway can increase the slot size and the  $TTL$  of each packet. As discussed in Section V-A, these parameters can be tuned to increase redundancy in the presence of poor links.

In our 52 node deployment, each slot was  $25ms$  with a  $TTL$  of 3. A maximum 802.15.4 packet takes  $4ms$  to transmit, meaning that each flooding cycle takes at most  $12ms$  of the slot time. Our TDMA cycle length was set to 64 slots such that each node would report once every 1.6 seconds. With additional generation capacity (through sources like solar) we






Fig. 10. GPS RSSI site survey with (a) gateway to forwarder coverage and (b) meter to forwarder coverage (c) final node placement with links

plan to eventually connect the remaining 450 homes in Les Anglais. We estimate this would require 25 forwarders with a TTL of 6 and a slot length of  $40ms$  which yields an update rate of approximately once per minute.

To ensure that Flash Flood is able to balance contention with load, we ran a simple experiment in a high-contention configuration (using Flash Flood version 1) to determine the worst-case limit of our forwarders. Figure 9 shows how the packet reception rate (PRR) is influenced by the number of forwarders. The bottom line shows the case where all forwarders are very close to each other in a high-density configuration that is the worst scenario for the capture effect. The top line shows the performance as the same nodes are spread apart 10 meters from each other. The source and gateway were out of range of each other and hence required forwarders to communicate. We see that even in the highly congested scenario, reception slowly increases up to 3 or 4 transmitters before it begins to fall off due to contention. In the more spread out scenario, PRR remains stable at more than 8 forwarders at a density higher than what we would see in practice. This indicates that with correct tiling of forwarders, even the simple Flash Flood protocol is quite effective.

### E. Deployment Methodology

Our deployment strategy was comprised of two main phases. First, we planned the physical network topology based on a GPS site survey that estimated device coverage. These ranges were used to place locations on satellite photos. Next, we installed the devices, collected data for approximately one month and then returned to iterate on the network topology.

Since our system has two classes of radio (high-powered forwarders and low-powered end-nodes), it was important to establish range estimates. We placed a gateway node set to continuously broadcast at the local operations facility. We then walked around with a forwarder mounted on top of a 15 ft pole that included GPS coordinates with its received signal values. Figure 10(a) shows a heat map of the signal strength around the town. Black areas denote regions where there was GPS but no gateway signal. Next, we repeated this experiment with the low-powered node mounted in the gateway location and the forwarder node being moved around at ground-level shown in Figure 10(b). This approximated the case where a meter needs to transmit to the nearest forwarder. We then used these

approximates to tile the town with five forwarders as shown in Figure 10(c). The dark lines indicate links used for forwarding traffic.

The installation process took approximately three weeks: one for installing distribution lines, one for installing service drops to homes and one for installing the meter hardware and forwarders. When installing the meters, we used a utility to configure the node channel, MAC address and AES encryption key, all of which are stored in EEPROM. After installing the system, we left each home in the off-state for one month while we monitored the network performance. As discussed in Section V-A the performance degraded significantly and required both additional forwarders as well as fine tuning of device placement to remain consistent.

## V. SYSTEM PERFORMANCE

As described in Section IV-D, our final deployment transmitted a packet once every 1.6 seconds from each of the 52 meters. At the time of writing this paper, the network has been operating smoothly for almost six months, and has been supplying power to customers for nearly four months.

### A. Network Performance

Immediately after our installation (before we energized the grid), we saw relatively good performance. As shown in Figure 11, the performance of the network significantly decreased over the first month of installation. Before activating the grid, we returned to Les Anglais and adjusted the positions of many of the SparkMeters as well as added an additional forwarder node in the southwest corner of the network. We also ran extensive tests with different TDMA parameters to rule out protocol malfunctions (for example packets arriving out of order). We determined that poorly performing nodes were solely a result of weak or intermittent signal strength. The final map in Figure 11 shows the network performance one week after the tuning. Since then, the network has remained stable.

The other critical metric to evaluate overall system performance is the node reporting latency. As described in Section IV, the reporting latency has a direct impact on the accuracy of billing and the overall quality of the user experience. Figure 12 shows the average and worst-case latency of packets received over a 24 hour period. Data was collected on the day

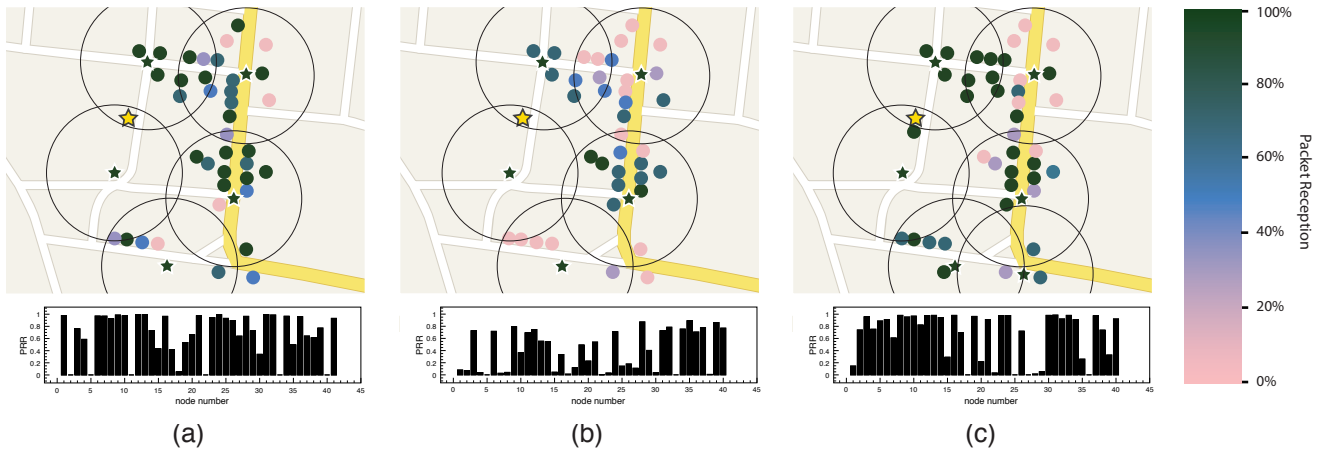


Fig. 11. PRR over a 1 hour window post-installation, 1 month later and then after tuning

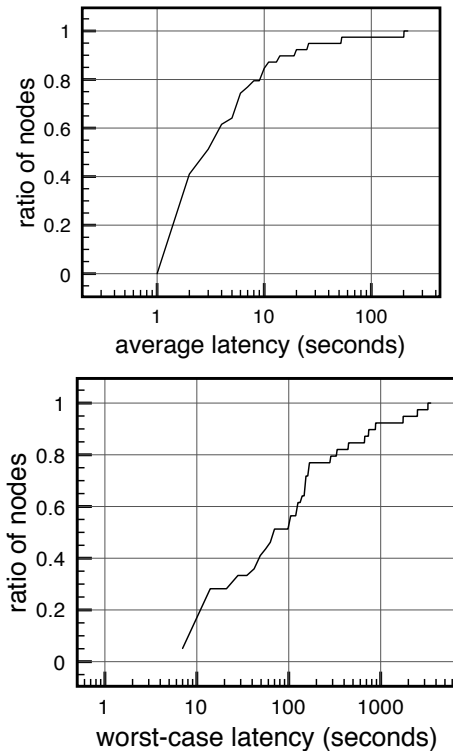


Fig. 12. CDF of average/worst meter latency

of a tropical storm to show the most extreme variations. We see that 40% of nodes have an average latency of 2 seconds or less and 80% of nodes have an average latency of 10 seconds or less. In the worst-case, 60% of nodes reply in less than 100 seconds with an overall worst case time of around 30 minutes which was likely due heavy rain.

### B. Data Fidelity

The most important aspect of our system is its ability to accurately meter generation sources and loads. Figure 13 compares traces from a single SparkMeter with high-resolution traces from the 3-phase meter. The traces were computed from

a 10-hour data set consisting of 11,000 points. The first set of traces (left) shows true power for both power phases as read by the 3-phase meter. At 2.3 hours, there is a significant dip in load caused by lights being shut off. The second set of traces (middle) compares voltage levels sensed by the 3-phase meter and an isolated SparkMeter. The last set of traces (right) compares frequency data between the 3-phase meter and a single SparkMeter. Traces indicate a high correlation in readings between the two meters and show the SparkMeter can be used to accurately estimate the state of the microgrid. In both sets, a shift in voltage and frequency occurs from 2.3 hours to 4 hours. This shift is caused by changes in load and are less pronounced in the SparkMeter because they are occurring on the opposite phase. In addition, oscillations in voltage and frequency are present after 6 hours. These are caused by the activation of the cell tower cooling system in response to increased ambient temperature.

Figure 14 shows pairwise voltage and frequency error spread for the 52 SparkMeters on the microgrid, sampled over a 10-hour window. The first plot (top) shows the bimodal nature of voltage values, a result of the fact that each SparkMeter serves one of two dedicated phases of generator output. Within each phase, 90% of samples fall within 2V of the average for that phase. The second plot (bottom) displays the error in frequency, which does not depend on the phase selected. Frequency values were found to be highly consistent, with 85% of samples falling within 0.5Hz of the average.

## VI. LESSONS LEARNED

Deploying microgrids in remote regions requires extensive planning and coordination of both personnel and supplies. We performed extensive functional testing in our lab before deploying the system, which paid off in that the system essentially performed as designed. What we could not test in the lab was how the system would respond under the volatile environmental conditions at the deployment site. Homes in Les Anglais have tin roofs which significantly attenuated the signal depending on where meters were mounted. Since it was difficult to energize the grid on a node-by-node basis, we had to install most of the nodes before being able to test connectivity. Having the option to operate nodes off of batteries (even if




Fig. 13. True power for two phases over a 10 hour period (left). Voltage measured at the generator (middle-top) as compared to an isolated SparkMeter (middle-bottom) along with frequency measured at the generator (right-top) compared to the SparkMeter (right-bottom).

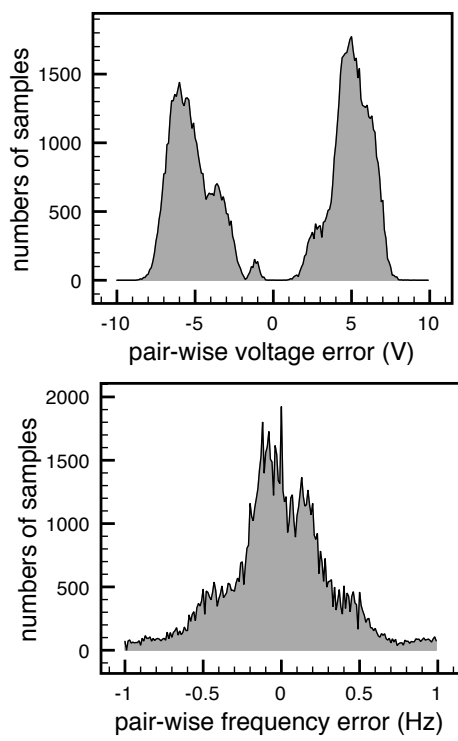


Fig. 14. Pairwise error of voltage and frequency across 52 SparkMeters

temporarily) would have saved time. We also saw that the RF characteristics changed as the network aged. We believe this change happened slowly, in which case careful trend monitoring could help preventatively isolate issues. Through this experience we also validated the power of visualizing geographically-based network information. Seeing node statistics in a simplified view on GPS coordinate maps became an invaluable debugging tool as the deployment progressed.

Another unanticipated outcome learned during our deployment was how automation impacts user interaction with the system. Our energy enforcement with quick reset and test has significantly increased per-user energy consumption compared to the previous manual circuit-breaker system. In the upgraded system, since enforcement happens on a minute-ly rather than daily basis, users are less discouraged from trying to consume closer to their limit. This has resulted in increasing the total load on the system by 500% of what we anticipated. The additional consumption increases the revenue potential while still bounding overall usage limits.

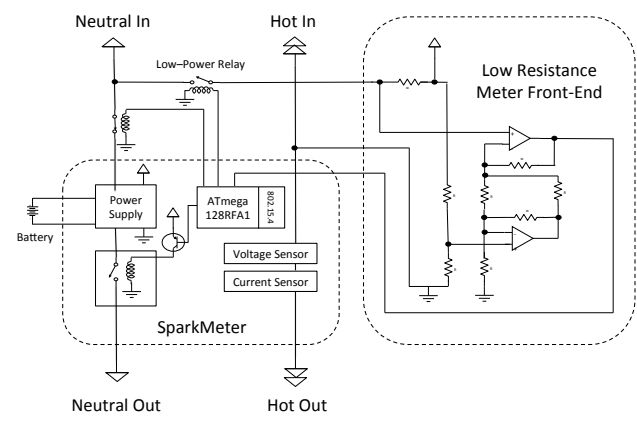


Fig. 15. Short-circuit measurement schematic

When we first powered-up the grid with all of the SparkMeters attached, the main breaker immediately tripped. It became clear that there was a short circuit somewhere in the system. Shorts can be traced with either an RF carrier transmitter like those used to trace circuits in homes or a sensitive multi-meter used to periodically probe resistance. RF transmitters often bleed past shorts providing inaccurate readings. Measuring resistance can be a tedious process if the network has multiple branches. It took our team nearly two full days of tracking down what turned out to be three faults in the wiring before we were able to eliminate all problems. In one case, the linemen climbing the pole to check for faulty wiring stepped on a bracket that caused another short. For this reason we propose updating our SparkMeter design to include functionality for tracing resistance in the event of a fault in the system.

Figure 15 shows a conceptual design of a circuit that would add low-value resistance measurement into the meters in the event of a power loss. The sensing front-end uses a set of relays to switch over to battery power during such an event. This is a major advantage of wireless as compared to PLC. Since meters inject voltage to determine resistance they cannot operate concurrently. One solution would be to leverage our TDMA protocol to avoid measurement collisions. Each meter can take a resistance measurement and then transmit the data back to the gateway for analysis.

Though beyond the scope of this paper, we believe there are techniques that can be used if approximate wire length estimates are known in order to localize microgrid shorts. Conversely, the circuit can also be used to trace open-circuits due to downed lines. Seeing a gradient map of resistances across the network would have drastically reduced our debugging time.



## VII. CONCLUSIONS AND FUTURE WORK

This paper showed the viability of using wireless sensor networking technology to manage microgrids in rural environments. As compared to solutions like power line communication, wireless offers advantages in terms of ease-of-deployment and robustness in the presence of faults. We discuss the performance of our deployment in Les Anglais, Haiti that includes a backend for fine-grained pre-payment of electricity along with a field-deployed metering and control infrastructure. This prototype system is currently powering 52 homes that no longer require expensive and toxic kerosene for lighting and are able to charge cellular phones and operate fans. In addition, it presents the opportunity for local storage of medicines that require refrigeration.

While this paper focuses on the hardware, software and networking components of the system, there are significant additional challenges related to the policy and economics of making systems like this viable. As future work, we intend to refine and scale the system such that it can be easily duplicated across regions that suffer from energy poverty. Scaling will require a management framework that can balance demand with multiple generation sources like solar panels and storage banks. We also need to design the hardware infrastructure in such a manner that it can easily be upgraded to higher power capacity metering systems to support future modernization of the infrastructure. As part of this effort, our next generation meters will include our proposed fault diagnostics circuit and will fit the standard enclosure fixtures used around the world. As the microgrid grows, the low-cost meters will be replaced with the latest variety of smart meter. Finally, we plan to evaluate different theft detection and management schemes using this deployment as a testbed.

## VIII. ACKNOWLEDGMENTS

This research was funded in part by the Link Foundation Energy Program, the Bertucci Graduate Fellowship program and TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. We would also like to thank the Carnegie Mellon Electricity Industry Center (CEIC) for their support.

## REFERENCES

- [1] I. E. Agency, "World energy outlook database," 2012.
- [2] R. Deshmukh, J. P. Carvalho, and A. Gambhir, "Sustainable development of renewable energy mini-grids for energy access: A framework for policy design," 2013.
- [3] G. of Haiti, "Haiti energy sector development plan 2007-2017," 2006.
- [4] M. Harper, "Review of strategies and technologies for demand-side management on isolated mini-grids," 2013.
- [5] M.-G. D. Manual, "Esmap technical paper 007," *Energy Sector Management Assistance Programme (ESMAP) UNDP/World Bank–April*, 2000.
- [6] T. Quetchenbach, M. Harper, J. Robinson IV, K. Hervin, N. Chase, C. Dorji, and A. Jacobson, "The gridshare solution: a smart grid approach to improve service provision on a renewable energy mini-grid in bhutan," *Environmental Research Letters*, vol. 8, no. 1, p. 014018, 2013.
- [7] J. Lu and K. Whitehouse, "Flash flooding: Exploiting the capture effect for rapid flooding in wireless sensor networks," in *Proceedings of IEEE INFOCOM, Rio de Janeiro*, 2009.
- [8] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele, "Low-power wireless bus," in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys '12. New York, NY, USA: ACM, 2012, pp. 1–14. [Online]. Available: <http://doi.acm.org/10.1145/2426656.2426658>
- [9] "DESI Power:<http://www.desipower.com/downloads/DESI-Power-Company-Profile.pdf>."
- [10] Chaudhuri, "Renewable energy in the sundarbans," *New Delhi: The Energy and Resources Institute*, 2009.
- [11] UNDP, "Case studies of sustainable development in practice: Triple wins for sustainable development." *United Nations Development Program*, 2012.
- [12] K. D. Kirubi C, Jacobson A and A. Mills., "Community-based electric micro-grids can contribute to rural development: Evidence from kenya." *World Development 37:12081221.*, 2009.
- [13] M. Khan and M. Iqbal., "Pre-feasibility study of stand-alone hybrid energy systems for applications in newfoundland," *Renewable Energy*, 2005.
- [14] S. S. Rana S., Chandra R. and M. Sodha, "Optimal mix of renewable energy resources to meet the electrical energy demand in villages of madhya pradesh." *Energy Conversion and Management*, 1998.
- [15] D. M. Kammen and C. E. Casillas, "The delivery of low-cost, low-carbon rural energy services," 2011.
- [16] "Devergy: <http://www.devergy.com> (viewed 10/10/2013)."
- [17] "Circuitor 2012a. Renewable Energies: <http://circuitor.com> (viewed 2/22/2013)."
- [18] "INENSUS: <http://www.inensus.com/en/products5.htm> (viewed 10/10/2013)."
- [19] "Power Hive: <http://powerhive.com> (viewed 10/10/2013)."
- [20] "GRAM Power: <http://www.grampower.com/> (viewed 10/10/2013)."
- [21] "Conlog <http://www.conlog.co.za/pages/ProductsServices/Single-Phase-Meters.html> (viewed 10/10/2013)."
- [22] Lifton, J., Feldmeier, M., Ono, Y., Lewis, C., Paradiso, J., "A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments," *International Conference on Information Processing in Sensor Networks (IPSN)*, April 2007.
- [23] Jiang X., Ly M V., Taneja J., Dutta P., and Culler D., "Experiences with a High-Fidelity Wireless Building Energy Auditing Network," *SenSys*, November 2009.
- [24] "<http://www.tendrilinc.com/> (viewed 2/20/2009)."
- [25] "<http://www.alertme.com/> (viewed 2/20/2009)."
- [26] "<http://www.trilliantinc.com> (viewed 2/20/2009)."
- [27] "<http://www.greenwavereality.com> (viewed 3/20/2010)."
- [28] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, "Efficient network flooding and time synchronization with glossy," in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, 2011, pp. 73–84.
- [29] M. Maróti, B. Kusy, G. Simon, and A. Lédeczi, "The flooding time synchronization protocol," in *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, ser. SenSys '04. New York, NY, USA: ACM, 2004, pp. 39–49. [Online]. Available: <http://doi.acm.org/10.1145/1031495.1031501>