

Neighborhood Watch: Security and Privacy Analysis of Automatic Meter Reading Systems

Ishtiaq Rouf*, Hossen Mustafa*,
Miao Xu, Wenyuan Xu†
University of South Carolina
{rouf, mustafah, xum, wyxu}@cec.sc.edu

Rob Miller
Applied Communication
Sciences
rob.d.miller@ieee.org

Marco Gruteser
Rutgers University
gruteser@winlab.rutgers.edu

ABSTRACT

Research on smart meters has shown that fine-grained energy usage data poses privacy risks since it allows inferences about activities inside homes. While smart meter deployments are very limited, more than 40 million meters in the United States have been equipped with Automatic Meter Reading (AMR) technology over the past decades. AMR utilizes wireless communication for remotely collecting usage data from electricity, gas, and water meters. Yet to the best of our knowledge, AMR has so far received no attention from the security research community. In this paper, we conduct a security and privacy analysis of this technology. Based on our reverse engineering and experimentation, we find that the technology lacks basic security measures to ensure privacy, integrity, and authenticity of the data. Moreover, the AMR meters we examined continuously broadcast their energy usage data over insecure wireless links every 30s, even though these broadcasts can only be received when a truck from the utility company passes by. We show how this design allows any individual to monitor energy usage from hundreds of homes in a neighborhood with modest technical effort and how this data allows identifying unoccupied residences or people's routines. To cope with the issues, we recommend security remedies, including a solution based on defensive jamming that may be easier to deploy than upgrading the meters themselves.

Categories and Subject Descriptors

C.2.0 [General]: Security and protection; C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Security, Experimentation

Keywords

AMR meter, Privacy, Spoofing, Reverse engineering

*Ishtiaq and Hossen contributed equally to this work.

†Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CCS'12, October 16–18, 2012, Raleigh, North Carolina, USA.
Copyright 2012 ACM 978-1-4503-1651-4/12/10 ...\$15.00.

1. INTRODUCTION

Much attention has been devoted to smart meters, as they play a central role in transforming the current electrical grid to the smart grid. The smart grid promises to improve the efficiency and reliability of existing grids by adding self-monitoring, self-diagnosis, demand-response, and communication capabilities. Privacy advocates have, however, immediately cautioned that fine-grained energy consumption data collected by smart meters could reveal sensitive personal information from homes, and have protested against the adoption of smart meters [1, 2, 3]. Researchers have already identified several specific types of information that can be derived from this energy data and pointed to associated privacy risks. For instance, burglars could monitor the power consumption of several households to identify temporarily vacant homes and time their break-ins [4]; a nosy landlord, employer, or even a stalker could estimate the number of residents in a household based on the frequency of power switches turned and the number of appliances simultaneously in use [4]; one could monitor the location of a resident inside the home based on the type of appliances being used [5]; health insurers could track eating, sleeping, and to some extent exercise habits by monitoring household appliance usage [6, 7]. Sufficiently fine-grained data even allows identifying the TV channel or movies being watched since television power consumption changes with the image being displayed [8]. Such snooping into personal residences is often viewed as particularly intrusive since the home is one of the last bastions of privacy. It is the locus that first gave rise to the concept of privacy laws [9] and a place where a particularly strong expectation of privacy still exists.

There are also security concerns that arise in smart meters, particularly about the integrity and authenticity of the reported data. Although smart meters have been designed to mitigate the long-standing problem of energy theft via physical tampering, their wireless module is known to present ample opportunity for dishonesty. For instance, it allows underreporting energy usage or inflating the utility bills of a neighbor [10, 7]. To ensure the trustworthiness of the meter data, NIST is developing guidelines for cryptography and key management strategies in smart meters [11]. Because of the on-going privacy debate and the yet-to-be-verified security strategies, smart meters have so far only undergone a few pilot deployments in the United States.

1.1 Automatic Meter Reading

There is, however, another enhanced meter technology that has already been widely deployed but has so far re-

ceived little attention. As of 2010, more than 47 million of these *Automatic Meter Reading* (AMR) systems were installed, representing more than one-third of the 144 million total U.S. residential, commercial, and industrial meters [12]. Once deployed, AMR meters will likely remain in operation for an extended period of time for several reasons. First, AMR systems enable utility companies to remotely collect consumption data of electricity, water, and gas — for example, with a receiver mounted on a drive-by truck. They therefore promise to reduce the cost of reading meters as well as reduce human errors in this process. Second, we have witnessed attempts to leverage the existing investments in AMR meters [13] to provide some of the functions of smart meters. For instance, existing *fixed network AMR* can report energy consumption data to both customers and providers in real time by connecting AMR meters to a network of radio repeaters and collectors [14], allowing utilities to better respond to demand changes (demand-response).

Motivated by this large existing deployment, in this paper we report a privacy and security analysis of a popular AMR meter system. Using a software radio platform we reverse engineered the wireless communication protocol and examined whether any of the above mentioned privacy and security risks associated with smart meters also exist in the AMR systems. We were also curious whether the considerable smart meter security research and public discussion have influenced the design of such systems.

1.2 Contributions and Findings

The step from traditional analog meters to AMR may seem like a minor technology upgrade compared to the envisioned smart meters, thus appearing unlikely to result in significant privacy and security issues. We found, however, that the risks are compounded due to the following reasons.

Unsecured Wireless Transmission. Smart meter research typically assumes that energy data is communicated to the utility over a secure channel to preserve the integrity of the readings and alleviate privacy risks [15, 11, 7, 16]. Thus, privacy risks of smart meters center on insider risks, misuse, and exploitation of the data at the utilities. The AMR meters that we studied, however, make data publicly available over unsecured wireless transmissions. They use a basic frequency hopping wireless communication protocol and show no evidence of attempting to ensure confidentiality, integrity, and authenticity of the data. The communication protocol can be reverse engineered with only a few days of effort and software radio equipment that is publicly available for about \$1,000 (GNU Radio with the Universal Software Radio Peripheral). We were able to both eavesdrop on messages as well as spoof messages to falsify the reading captured by a commonly used ‘walk-by’ reader.

Continuous Broadcast of Fine-Grained Energy Data.

The meter we examined continuously broadcasts its energy consumption, even if no receiver is present. Approximately once a month, the meter is being read by a utility truck that drives by. However, the meter simply broadcasts its reading every 30s around the clock. We also found the communication range of AMR meters to be larger than expected. Packets from gas meters and electric meters can be successfully received from up to 70m and 300m, respectively, using a generic 5 dBi antenna and an off-the-shelf low noise amplifier (LNA). In the neighborhood where we tested, we were able to receive packets from 106 electric meters using a basic

antenna and 485 meters by adding an LNA at a single meter location.

Neighborhood Monitoring. AMR meters make it possible for anybody with sufficient technical skills to monitor real-time energy consumption patterns in an entire neighborhood. We built a live RF sniffer that collects energy consumption records by eavesdropping on the periodic wireless packets. Compared with the more fine-grained data obtained from direct visual observation, we found that the data obtained by RF sniffing is still sufficient to identify the same appliance usage events. It can therefore be helpful in inferring residents’ daily routines. This is particularly concerning because wireless eavesdropping facilitates the monitoring of hundreds of residences in neighborhood from a single location with a lower risk of detection than direct visual observation of the residences.

Defenses for Legacy AMR Meters. We recommended several remedies to alleviate these risks on legacy meters. They range from policies such as occasional manual cross-checking of the readings to ensure integrity, over meter upgrades with cryptographic protocols, to a defensive jamming solution that can be implemented by adding a simple hardware component next to an existing meter.

The rest of the paper is organized as follows. We present a background overview of AMR in Section 2. In Section 3, we describe our reverse engineering endeavor to discover details of the proprietary communication protocol, and show spoofing attacks. We then reveal our finding on using AMR meters to monitor energy usage in a neighborhood in Section 4 and to identify people’s routines in Section 5. In Section 6, we recommend security remedies. Finally, we survey the related work in Section 7 and conclude the paper in Section 8.

2. BACKGROUND

Automatic Meter Reading (AMR) is a technology that autonomously collects the consumption and status data from utility meters (e.g., electric, gas, or water meters) and delivers the data to utility providers for billing or analysis purposes. The concept of AMR was proposed in the 1960s [17], and the first AMR design was documented in a patent by Paraskevagos in 1972 [18]. This early version of an AMR system used telephone lines to automatically transmit meter readings to a remote receiver. Later versions of AMR adopted power line communication, low power radio frequency (RF) communication, satellite communication, etc. Among these technologies, RF communication is the most cost-effective solution and has been widely used in residential AMR systems. In this paper, we focus on AMR systems that utilize RF communication.

2.1 AMR Architecture

AMR systems consist of two main components: (1) *AMR Meters* that collect and transmit consumption data, and (2) *AMR Readers* that receive and forward the consumption data sent by meters to a central collection point for billing, diagnosis, and analysis.

AMR Meters. AMR meters (hereafter *meters*) measure the total consumption of electricity, gas, or water. Regardless of what meters are measuring, their core components remain the same. Each RF-based meter is comprised of a metering engine and an Encoder-Receiver-Transmitter (ERT). The metering engine measures the consumption through a



Figure 1: Generic AMR meters. [Left to right] A stand-alone gas meter, a gas meter inside gas flow measuring chambers, and an electric meter.

mechanical dial that rotates at a speed proportional to the amount of consumption. With the help of electromechanical or electro-optical interfaces, the movements of dials are converted into digital numbers. The ERT consists of a microprocessor and a low-power radio transmitter. It processes the meter reading and periodically reports information such as meter ID, meter reading, tamper status, etc.

Depending on what meters are measuring, their appearances, communication protocols, and power supplies can differ. Fig. 1 shows two representative residential gas meters and an electric meter. Electric meters are conveniently powered by the main electricity supply line, while gas and water meters operate on sealed batteries designed to last up to 20 years [14]. The battery constraints of gas and water meters usually lead to longer intervals between energy reports. We analyze both electric and gas meters with an emphasis of electric meters.

AMR Readers. To capture the meter readings and relay them to a central collection point, one or more AMR readers (hereafter *readers*) are required. Readers interpret the signals and deliver the meter IDs along with other information to a central collection point. Three categories of readers are used in the utility industry: (1) handheld devices for field investigation or walk-by meter reading, (2) highly sensitive mobile collectors for drive-by meter reading, and (3) a network of permanently installed collectors and repeaters for reporting AMR meter readings in real time, (aka. *fixed network AMR*) [19].

Both handheld devices and mobile collectors require personnel to walk or drive by locations where the meters are installed, and total utility consumption can only be updated as frequently as the walk-by or drive-by events occur. A fixed network AMR system requires higher infrastructure investment, but does not need delegated drivers or ‘walkers’ for data collecting, and can provide continuous energy consumption updates to the utility.

Since we were unable to get full access to mobile collectors or fixed network collectors, we show our findings using a handheld collector. Because the main function of all three types of AMR readers is to collect meter readings, we believe that our findings provide insight for the other types of AMR readers.

2.2 AMR Communication Protocol

The communication protocol between meters and readers is proprietary. Even so, a survey of information from supplier websites and patents [20] provides a rough idea about the communication protocol, with some information proving

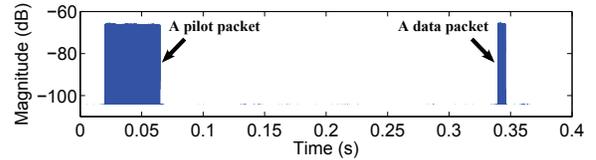


Figure 2: An AMR transmission is comprised of two packets. A pilot packet is transmitted approximately 275 ms before a data packet that contains the actual meter status update.

to be inapplicable to the models of meters that we studied and some proving to be pertinent. We learned that most meters operate in the 915-MHz ISM band, use simple modulation schemes such as on-off keying (OOK) or frequency shift keying (FSK), and incorporate Manchester encoding schemes. To avoid packet collision, meters implement frequency hopping, where packet transmissions repeatedly cycle through a pre-determined sequence of channels. Each packet contains the meter ID, reading, device type, and tampering status.

AMR systems support two types of communication models: ‘wake-up’ and ‘bubble-up’. Wake-up systems use two-way communication, whereby a reader transmits an activation signal to wake up and interrogate one or more meters. Bubble-up models use one-way communication, whereby meters periodically broadcast the meter readings. Wake-up models are primarily used in battery-operated gas and water meters, while bubble-up models are used mainly in electric meters [14]. Interestingly, we discovered that the gas meter that we investigated also works in bubble-up fashion.

3. SECURITY ANALYSIS OF AMR METERS

Besides the effort of detecting physical meter tampering, we have found no evidence that security was considered during the AMR meter design. Since the wide deployment of AMR meters, there has been sporadic exploration into system characteristics [21, 10]. However, none of the previous work has taken a comprehensive look at the deployed system to determine how an attacker might misuse it.

In this section, we investigate the following issues: (1) How easy is it to reverse-engineer the communication protocol? (2) Are spoofing attacks possible?

3.1 Equipment

The primary purpose of our work is to raise awareness about oft-neglected areas, not to encourage misuse; hence, we have refrained from disclosing details of the meters being studied. For our study, we used the equipment from the following three categories.

AMR Meters. We selected electric and gas meters that have been widely deployed throughout the United States. In addition to meters installed in our neighborhood, we acquired second-hand electric and gas meters to conduct experiments both in the lab and outdoor.

Meter Readers. We obtained a generic handheld AMR ERT module reader used by meter inspectors for field interrogation. This handheld AMR reader works with the selected meters and can read the meter ID, meter reading, meter type, and physical tamper status. In our experiments, we primarily used it for interpreting meter packets. We did manage to briefly access (1 hour) advanced AMR collectors used by utility companies to test our spoofing attacks.

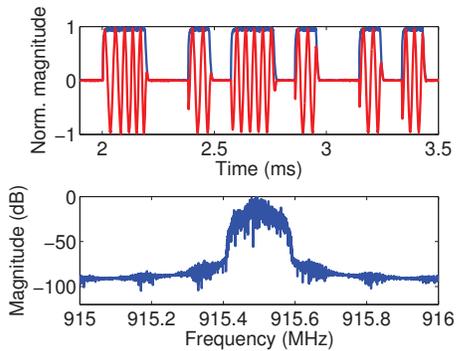


Figure 3: A captured AMR transmission. The top plot shows the signal time series (both magnitude and in-phase), while the bottom plot displays its spectrum. The plots suggest the use of OOK modulation and Manchester encoding.

However, we did not report the experiment results due to the mutual agreement with the utility company.

Raw Signal Sniffer. We were interested in performing analysis using low-cost, off-the-shelf equipment; hence, we did not use professional signal processing instruments such as Vector Signal Analyzers or Spectrum Analyzers. Instead, we used a low-cost software defined radio platform, i.e., the Universal Software Radio Peripheral (USRP) [22] to capture raw AMR signals. We primarily used the first generation USRP with limited instantaneous bandwidth capability, and occasionally tested on the second generation that can monitor a wider frequency band. The daughterboards used in our experiment include an RFX900 and a WBX daughterboard. Both daughterboards cover the frequency range of AMR meters, which is centered at 915 MHz.

For ease of reading, from this point onward we refer to the electric meter as *meter-E*, the gas meter as *meter-G*, and the handheld reader as *ERT reader*.

3.2 Reverse Engineering AMR Communication Protocols

Without insider information, we rely on reverse-engineering to discover the meters’ communication protocol details.

Capturing the First Packet. The first step of reverse engineering is to capture a few transmissions from each meter. Surprisingly, this step turned out to be very difficult for several reasons. First, it took a long time to detect and capture AMR activities. Both meter-E and meter-G work in the bubble-up model, where they periodically broadcast packets and do not respond to any activation signals. Second, it is difficult to capture a ‘clear’ transmission from meters without the specific channel and other physical layer information. Meters hop through a frequency band larger than what our signal sniffers can monitor. During the exploration phase, we used the first generation USRP that is able to monitor at most 8 MHz (with 16-bit I/Q samples), a fraction of what meters cover. Due to this limitation, we captured highly distorted packets which were beyond decoding. Moreover, several other electronic devices operate in the 915-MHz ISM band, and thus resulted in interference and confusion. As an example, we observed wireless transmissions from cordless phones.

To streamline our reverse-engineering process, we built detection software to capture, replay, and verify signals. First,

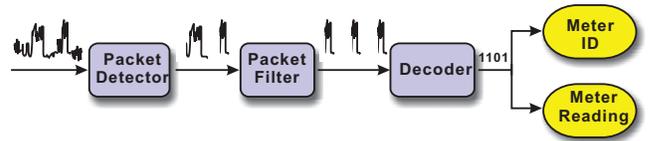


Figure 4: The flow chart of our live decoder.

we developed an activity detector to automatically record signals that are above the ambient noise floor. Then, candidate signals were extracted and verified by retransmitting them with the USRP. The signals that induced proper display on the ERT reader were the ones transmitted by AMR meters. In the end, we discovered that each transmission from AMR meters consists of a pilot packet with a length of 44.9 ms, and a data packet of 5.8 ms. Both packets were transmitted on the same channel with separation of 275 ms as shown in Fig. 2. Interestingly, we found that the ERT reader does not need the pilot packet to correctly interpret the data packet.

Decoding Packets. Lacking an instrument to interpret the pilot packets, we focused on decoding the data packets. To convert raw signals to bit streams and then to meaningful fields, we needed to identify the modulation scheme, baud rate, encoding scheme, and packet format. After a quick examination of the signals in the time and frequency domains (as shown in Fig. 3), we confirmed that meters use on-off keying (OOK) as their modulation scheme and an inverted version of Manchester encoding scheme.

Additionally, we discovered that the baud rate of meters is 16.5 kBd, and the duration of the high level of each bit has to be slightly longer than the duration of the low level. Packets with the improper ratio of high and low level durations will be considered invalid and are ignored by the ERT reader. It is unclear what the intention was for the different durations of high and low levels, but this difference did induce extra effort during our reverse engineering endeavor.

Successful demodulation and decoding returned a stream of 96 bits in each packet. With the help of the ERT reader, we were able to obtain several pairs of bit streams and their corresponding meter IDs, meter readings, tamper status, etc. Differential analysis over bit streams revealed the packet format, which contains a 24-bit meter ID, a 22-bit meter reading, and a 16-bit CRC checksum. We found that both meter-E and meter-G use the same packet format.

Characterizing Channels. To find details about the transmission channels of meters, we scanned through the entire 915-MHz ISM band (i.e., 902-928 MHz). We found a cluster of channels that are used by meters and made two interesting observations during this phase of experimentation: (1) Channels are separated by multiples of 200 kHz and not all candidate channels in the 915-MHz ISM band are used. (2) Meter-E cycles through a sequence of 50 pre-determined channels every 25 minutes.

Building a Live Decoder. After identifying details of the meter communication protocol, we developed a live decoder that monitors channel activities and outputs the meter ID and meter reading immediately after a packet is received. The live decoder consists of Python scripts that utilize signal processing libraries in GNU Radio. As shown in Fig. 4, the live decoder continuously samples the channels around 915 MHz at a rate of 4 MHz. We note that in such a setup, only packets transmitted in the range of 913 MHz to 917



Figure 5: Spoofing attack validation: The LED display of the ERT reader received the spoofed packet with an ID of 11223344 and data reading of 1234.

MHz can be captured with little distortion and can survive decoding. We will discuss our effort in capturing packets transmitted in other frequency ranges in Section 4.

Once the packet detector identifies high energy in the channel, it extracts the complete packet and passes the sampled data to a packet filter. To filter out non-AMR packets and pilot packets, we first discard any packet whose length mismatches with that of expected data packets. Next, we perform histogram analysis to discard distorted packets. The underlying observation is that, ideally, the amplitudes of a Manchester-encoded signal should cluster around two sets: one set mapping to the low level and the other mapping to the high level, as illustrated by Fig. 2. Any packet whose amplitude spectrum is not evenly divided into high and low levels is likely to be a distorted packet. Finally, ‘clear’ meter packets are passed to the decoder for extracting meter ID and meter reading.

Lessons Learned. At the end of the reverse engineering process, we came to the following conclusions.

- *Reverse Engineering requires modest effort.* With a communication and computer engineering background, one can reverse engineer the meter communication protocol with reasonable effort using off-the-shelf equipment (an ERT reader and USRP with an RFX900) costing \$1000 at the time of our experiments.
- *No Encryption.* No encryption algorithms are used, which makes it possible for *anyone* to eavesdrop on the real time consumption of customers with ‘bubble-up’ meters. For customers with ‘wake-up’ meters, it is foreseeable that their consumption data can be eavesdropped on at arbitrary rates using activation signals, since those signals are also not protected by cryptographic mechanisms.
- *Battery Drain Attacks.* After receiving an activation signal, ‘wake-up’ meters will immediately transmit a packet. Thus, they are vulnerable to battery drain attacks.

3.3 Packet Spoofing

After AMR meters have been installed, most customers and utilities trust the integrity of the collected meter readings, since AMR meters reduce human errors associated with the traditional analog meter collecting process. However, such a trust relationship must be questioned, should a malicious attacker be able to forge packets containing arbitrary data and successfully deliver them into the provider’s data collectors. Thus, following our successful reverse engineering step, we examined the feasibility of launching spoofing attacks.

To transmit a spoofed packet with an arbitrary meter ID and meter reading, we generated a properly formulated packet using Manchester encoding and OOK modulation.

Then, the fake AMR data packet was up-converted and transmitted at one of the channels used by meters. We have tested our spoofing attacks on the following three monitoring devices with gradually improved complexity: (1) a generic handheld collector (the ERT reader), (2) a more advanced data collector commonly used by field investigators in utility companies, and (3) a sophisticated mobile collector used by utility companies to gather meter readings from a vehicle driven in a fixed route periodically.

The authors and utility company agree that disclosing the results conducted at the utility company would not enhance their systems. Therefore, we exclude the test results obtained on the advanced handheld data collector and the sophisticated mobile collector, and only reveal our findings using the ERT reader.

Observations. Our experiments using the ERT reader reveal the following findings.

- *No Authentication.* The ERT reader accepts any AMR transmission with a proper packet format. Fig. 5 shows that the ERT reader accepted a spoofed packet with information of our choice: meter ID of 11223344 and meter reading of 1234.
- *No Input Validation.* When receiving multiple packets with the same meter ID but conflicting meter readings, the ERT reader will accept the packet with the strongest signal without reporting any warning. We note that even if a meter collector is sophisticated enough to keep track of all received packets for a conflicting test, an adversary can easily jam and block packets sent by a legitimate meter and let the meter collector only receive her spoofed packets.

4. NEIGHBORHOOD MONITORING

AMR meters pose immediate privacy risks as they broadcast meter readings in plaintext. Each packet contains a meter ID and meter reading. Given a specific household address, it is usually trivial to identify the associated meter ID because meters tend to be installed in publicly accessible locations (e.g., exterior walls of residential houses), and the meter ID is printed on the front face of meters. In this section, we explore whether an adversary can monitor a larger number of homes in a neighborhood simultaneously. This depends on the range of the transmission and propagation loss. Our software radio eavesdropping approach is also complicated by the frequency hopping feature, since it cannot monitor the entire set of frequencies simultaneously.

Experiment Setup. Unless specified, our basic eavesdropping experiments were conducted from inside the apartment of one of the authors. The antenna dedicated to eavesdropping was mounted against a third-floor window overlooking a slope with several buildings. All meters in this neighborhood are the same type as meter-E, and each meter transmits 1 packet every 30 seconds. For the majority of our eavesdropping experiments, we used one USRP (the first generation), with an RFX900 and a 5 dBi omnidirectional antenna, mimicking a narrowband receiver that can monitor a fraction of all channels (e.g., 17 channels centered at 915 MHz).

Eavesdropping Range. The first task was to estimate the eavesdropping range against meter-E using a basic 5 dBi dipole antenna in a few real world environments. We tested two representative locations in an author’s state: a rural

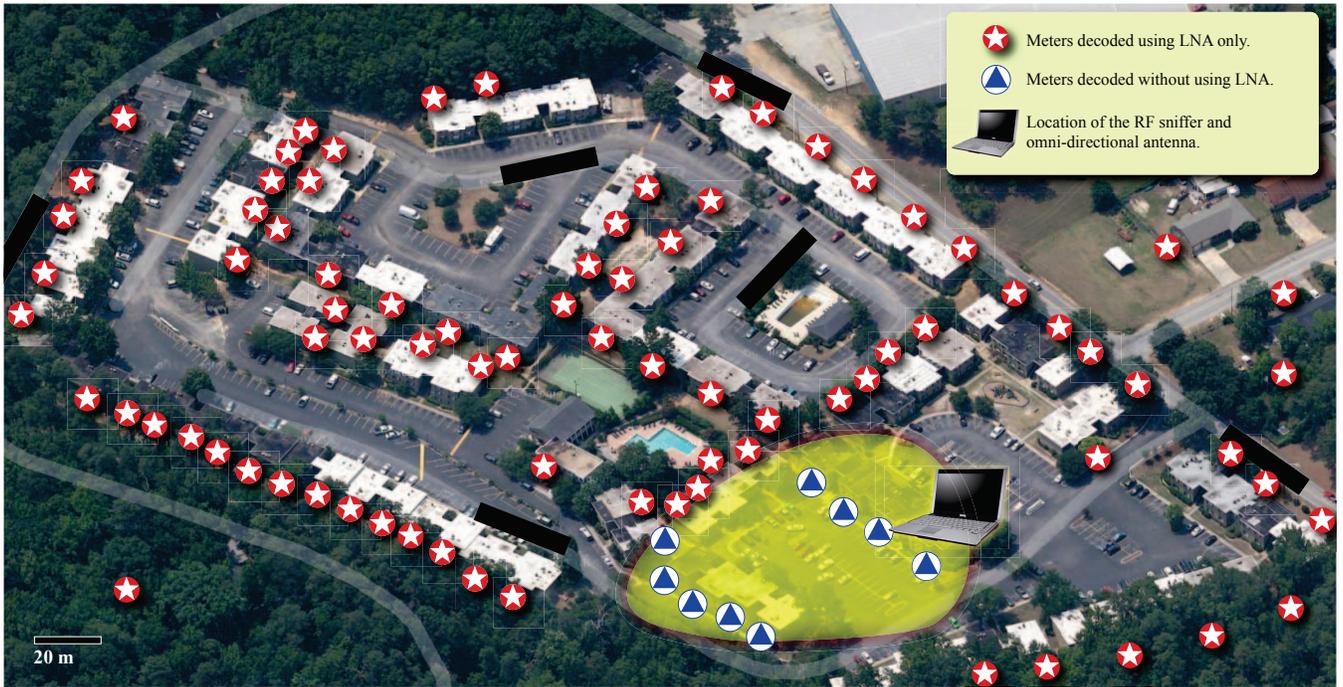


Figure 6: An aerial view of the neighborhood where we performed our eavesdropping experiments. Each blue triangle or red star represents a group of four or five meters mounted in a cluster on an exterior wall. Using an LNA and a 5 dBi omnidirectional antenna, we were able to monitor all meters in the neighborhood. Some sniffed meters may be out of the scope of this view.

area with sparse two-story independent houses and an urban area with several connected three-story apartment buildings (e.g., our basic experiment location). Since we were interested in determining the physical range of our eavesdropper, we considered a meter within eavesdropping range if at least *one* of its packets was intercepted and decoded successfully over the entire listening duration (1 to 4 hours).

We were able to decode packets from as far as 150m away in the rural area, and up to 70m in the urban area. We believe that the range difference is caused by the terrain variance. The rural area has far fewer obstacles (e.g., buildings) to hamper radio propagation (e.g., fading and multipath effects). Although the eavesdropping range in other environments may differ, our results indicate that an attacker should be able to sniff packets in any environment without entering private property.

Boosted Eavesdropping Range. To boost the range at low cost, we added a commercially available low noise amplifier (LNA) [23] to the antenna. The LNA provided 21 dB gain, and increased the eavesdropping range in the urban area from 70m to more than 300m for meter-E and from 15m to 70m for meter-G, as summarized in Tab. 1. Meter-G has a smaller range because it is battery-powered and transmits at a lower power level.

Fig. 6 provides an aerial view of the physical range of eavesdropping and the terrain variation. The laptop icon denotes the location of the eavesdropper. Without an LNA, the eavesdropper can decode packets sent by meters located at blue triangles. Once the LNA was added to the basic setup, we were able to collect data from a larger number of meters, denoted by the red stars. The underlying principle of increasing the receiving range is that an LNA am-

plifies the received signal strength (RSS) of each packet and thus increases the likelihood of successful decoding. To illustrate, Fig. 7 depicts the RSS of one meter located 15m from the eavesdropper when an LNA was and was not used. The usage of an LNA boosted the receiving range by several multiples, which enabled us to monitor meters further away. Granted, there are other ways to boost the eavesdropping range, but our intention was to show that the eavesdropping range can be increased using inexpensive hardware.

The Number of Observed Meters. To measure the total number of observed meters, we utilized two RF sniffers: a narrowband sniffer monitoring a 4 MHz frequency band and a wideband sniffer monitoring 12.5 MHz. The narrowband sniffer received packets from 72 meters *without* the LNA and 161 meters *with* an LNA. The wideband sniffer could receive 106 meters *without* an LNA and 485 *with* an LNA, which is more than the total number of apartments in the neighborhood (408 units). We believe some of the observed meters are located in the nearby region.

Increasing Packet Reception Rate. We use the number of received packets per hour (pph) to evaluate packet reception rate. A larger pph maps to a more frequent update on customer energy consumption and a high level of information leakage.

We observed that an LNA does help to boost the eavesdropping range and the number of observed meters, but it

Range	w/o LNA	w LNA
Meter-E	70m	300m
Meter-G	15m	70m

Table 1: Eavesdropping range for a gas and electric meters with and without an LNA.

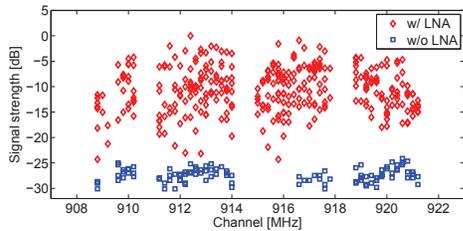


Figure 7: An illustration of boosting RSS using an LNA. We monitored packets sent by a meter 15m from the receiver.

also reduces the packet reception rates of the meters that could be heard without an LNA. Ideally, a narrowband sniffer that monitors 4 MHz centered at 915 MHz can hear about 30% of all transmission (40 pph), since meters tend to transmit around 915 MHz. Without an LNA we could receive 6.65 pph per meter on average and 27 pph maximum, while with an LNA, we could only receive 3.96 pph per meter on average and 27 pph maximum. The addition of the LNA undoubtedly increases the co-channel interference (CCI). Hence, our proof-of-concept sniffer sees more packet collisions.

In addition, a wideband sniffer can slightly improve both the average pph (7.03 pph) and the maximum ones (30 pph). The less-than-expected improvement is probably because the wider the receiving frequency band is, the more likely concurrent transmissions in different channels collide. Detailed distribution of pph for all meters is depicted in Fig. 8. We will show in the later section that even at a low reception rate, it is feasible to identify sensitive information of the residents, such as whether the residents are at home.

A few methods can be used to increase received pph. (1) A sophisticated decoding scheme. For instance, utilizing capturing effects, we can at least decode the strongest packet among collided packets, if the RSS of the strongest one is larger than other packets by a threshold factor [24]. Further, advanced signal processing techniques such as successive interference cancellation and multi-user detection can be combined with multi-antenna techniques [25] (e.g., beamforming and space-time adaptive processing [25]). (2) Monitoring the entire frequency range. A platform monitoring the entire frequency range can be used to capture packets transmitted at all channels. However, one would need a very powerful computer to process data at the rate and/or significant algorithm refinement to decode concurrent transmissions at different channels, which are outside the scope of our effort. (3) Dedicating one RF sniffer to monitor one meter. A narrowband RF sniffer can hop through the same channel sequence as the target meter to receive packets. Our experiments show that such a sniffer could achieve 88.5 pph monitoring a meter that is 10m away without an LNA.

Neighborhood Monitoring. Wireless monitoring allows the gathering of meter readings in an inconspicuous manner from a larger number of homes. By RF eavesdropping using a cheap antenna and a low-cost LNA, we were able to obtain an hourly distribution of power consumption in the authors’ neighborhood, as shown in Fig. 9. Since the precision of the wireless meter readings is 160Wh, each bar in the figure represents 160Wh more consumption than the one to its immediate left.

Consider, for example, that about 27 meters consumed less than 160Wh per hour on average, indicating the corre-

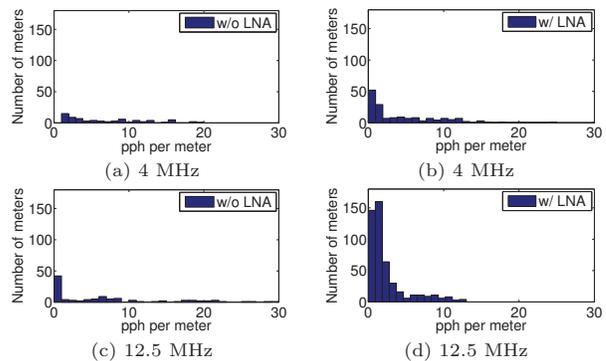


Figure 8: Histogram of received packets per hour (pph) from each meter using a narrowband sniffer (4 MHz) or a wideband sniffer (12.5 MHz).

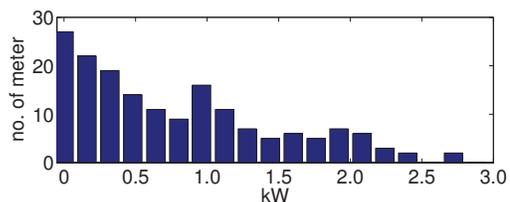


Figure 9: The distribution of electricity consumption for meters in the author’s neighborhood. 27 meters exhibited less than 160Wh hourly power consumption, indicating 27 apartments unoccupied.

sponding apartment units were likely to be unoccupied at the time of our experiments. This is an example of potentially sensitive information that can easily be obtained on this neighborhood scale. In this experiment, we were only able to receive a few packets per hour (pph) for a large portion of meters. Methods to increase received packet rates are available and therefore finer granularity data and additional sensitive information from the neighborhood could likely be obtained. We will examine this next.

5. INFERRING HOUSEHOLD EVENTS

We now study to what extent it is possible to infer detailed household activities and events from the obtained data—are the risks similar to those of smart meters? The lower update rate of 30s and high packet loss in neighborhood monitoring may suggest that this is more difficult.

Since no smart meters with fine-grained data are available in this neighborhood, we address this question by comparing our data with fine-grained data obtained from direct visual observation of the meter. To reduce the manual labor required in this process, we have implemented two automated visual observation mechanisms of a meter’s on-board LCD display and infrared (IR) LED using cameras or IR photodiodes, respectively. We considered them as the baseline schemes for comparing the level of privacy risks caused by wireless sniffing¹.

5.1 Automated LCD Screen Monitoring

The LCD shows the accumulated meter reading in digits with a resolution of 1 kiloWatt-hour (kWh) and the rate of consumption by a few ‘dots’, which are displayed on the

¹We summarized the highlights of the comparison among these three methods in Tab. 2

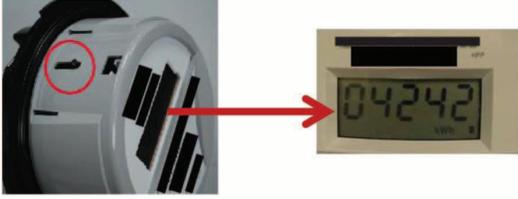


Figure 10: Meter-E disburses the energy consumption via an infrared LED highlighted by the red circle and the LCD display.

corner of the LCD, as shown in Fig. 10. The set of dots are turned on and off in a sequence such that they appear as if a digital wheel is spinning. Once one Watt-hour (Wh) is consumed, one of the dots toggles. To capture every 1 Wh consumption, our camera-based monitoring system tracks the changes of the dots on an LCD screen. Our system consists of a wireless network camera (AXIS 207W Network Camera) filming the LCD display and a laptop for data processing. Once the video is streamed to the laptop, the processing algorithm locates the area of dots, identifies every dot toggle, and generates an electricity consumption trace.

In our experiments, we mounted the network camera 0.3m from a meter to record its LCD display and set the camera at its 10 frames per second (fps). This rate ensures the recording of every dot flash unless a household consumes electricity at a rate higher than 18 kWh. Although our camera has to be located no more than 0.5m from the meter because of its low resolution (640×480), in practice, with a higher resolution, the camera can be mounted at a hidden location further from the meter.

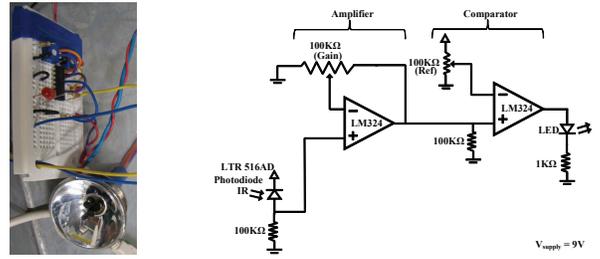
5.2 Infrared LED Monitoring

The IR LED on an electric meter flashes every time 1 Wh is consumed. To capture IR flashes, we designed an IR sensing circuit (as shown in Fig. 11), which works as follows. Without an IR signal to stimulate the IR photodiode, the output voltage of the low power amplifier (LM324) of the comparator will be low. Once the IR photodiode detects an IR light, the output voltage becomes high (5V). We record the output of the amplifier using a National Instruments USB-6009 data acquisition unit (DAQ) at a rate of 1 kHz, which is sufficient to capture IR flashes. In practice, the IR sensor can be combined with a microprocessor and an RF module to sense and report IR flashes wirelessly.

5.3 Experiments and Results

Here, we examine whether RF sniffing can reveal information about (1) *Daily routines*, like a resident’s sleeping cycle or work shift information. (2) *Appliance activities*, like the usage patterns of appliances in a household (e.g., when an appliance is on or off).

To evaluate whether RF sniffing can reveal daily routines and appliance activities, we measured two metrics: (1) *Step changes*, which are ‘jumps’ between two consecutive power consumption levels [26]. Let $\Phi(t)$ be the total power consumption rate (Watt) at time t . A step change happens when $|\Phi(t + \Delta t) - \Phi(t)| > \alpha$, where α is a threshold value determined by measurement noise. The number of step changes is an effective metric for detecting customers’ daily routines [6]. (2) *Activities*, which are appliances being turned on or off. To identify the number of step changes



(a) Detection circuit (b) Circuit diagram

Figure 11: IR LED flash detection circuit.

and to recognize appliance status changes, we developed our algorithms based upon Hart’s algorithm [26], which utilizes edge detection to identify step changes and recognizes appliance activities by matching step changes with appliance power consumption signatures.

Feasibility of Inferring Sensitive Information. RF eavesdropping suffers from low granularity of obtained data, since a meter-E broadcasts its reading every 30 seconds with a data precision of 160 Wh, and unpredictable channel environments and frequency hopping make eavesdropping unreliable. A key question to answer is whether RF eavesdropping can obtain data that suffice to infer sensitive information, a concern that arose in smart meters.

To answer the question, we conducted two experiments. In the first experiment, we monitored a meter of an apartment with one resident for 24 hours in late spring when heating was not used. The meter was installed in a private room with lighting, and thus the camera captured dot changes throughout the entire experiment. As expected, both camera and IR-based methods captured data with higher granularity than RF sniffing. To evaluate RF sniffing with various levels of receiving capability, we emulated the captured consumption data at multiple packet receiving rates (i.e., pph). We observed that when an attacker can eavesdrop a reasonable percentage of packets (e.g., 25 pph), enough *step changes* can be captured to identify high-load appliance activities (e.g., water heater turned on). Even when the received packet rate is as low as 2 pph, daily routines still can be inferred (Fig. 12).

The second experiment occurred in July when air conditioners were used to cope with the summer heat. We monitored the energy consumption using an RF sniffer in one author’s neighborhood. Figure 13 shows the electricity consumptions of one of the neighbors, from which we observe a pattern: The owners left for work on weekdays and stayed at home over weekends. Furthermore, we can infer the daily routines easily: The owner got up at 7 am, left for work around 9 am, and returned home around 6 pm on Friday.

In conclusion, AMR meters allow similar sensitive inferences as smart meters. What’s worse is that AMR enables

Method	Camera	IR Photodiode	RF Sniffer
Granularity	1 Wh	1 Wh	160 Wh
Range	~ 0.1-10m	~ 0.01m	~ 300m
Multiple meters	✓	×	✓
Light sensitive	✓	×	×
Line of sight	✓	✓	×
Daily routines	✓	✓	✓
Appliance activities	✓	✓	×

Table 2: Comparison of three monitoring methods.

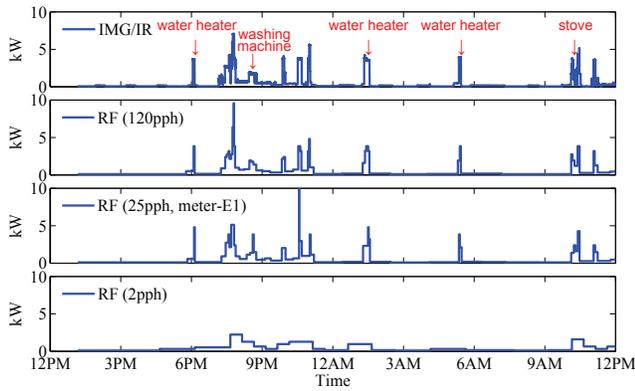


Figure 12: An RF sniffer can collect data that suffice to infer daily routines, which could be misused by thieves. [Top to bottom] The electricity consumption over a 24-hour period that are collected using (1) a camera or an IR photodiode circuitry; (2) an ideal RF sniffer receiving all packets; (3) a real meter being studied; (4) a narrowband RF sniffer that eavesdrops on one channel only.

any one, not just utility companies, to obtain sensitive information.

6. DEFENSE STRATEGIES

Automatic Meter Readers are vulnerable to spoofing attacks and privacy breaches because packets are sent in plaintext. We discuss a few strategies to improve the security and privacy of meters. The strongest level of protection would require a redesign of the communication protocol as outlined in Section 6.2. There are, however, possible jamming-based defenses for legacy meters that can raise the bar for attacks, and can be deployed more rapidly at a lower cost.

6.1 Spoofing Defenses for Legacy Meters

A few strategies are available to mitigate RF spoofing attacks for deployed meters without modifying the meters. The first one is radio fingerprinting techniques, which can differentiate amongst transmitters (e.g., real meters or attackers in this case) by exploiting device levels imperfections [27] or unique channel responses [28]. Secondly, anomaly detection over a collection of meter readings can identify a sudden usage change and raise an alarm to perform a spoofing investigation. Furthermore, utility personnel can check the meter reading in person occasionally to detect spoofing attacks.

6.2 Cryptographic Mechanisms

A complete solution would use cryptographic mechanisms to achieve authenticity, integrity, and confidentiality. For instance, the data packets can be encrypted using standard block encryption algorithms and augmented with a digital signature for authentication. As such, an attacker cannot casually eavesdrop the wireless communication and obtain sensitive power consumption data of consumers. Without the private key of the meter, the attacker cannot forge the signature of meters and claim arbitrary meter readings. It would also be a good practice to transmit a meter reading only when needed. For example, letting a drive-by reader wake up AMR meters appears more appropriate.

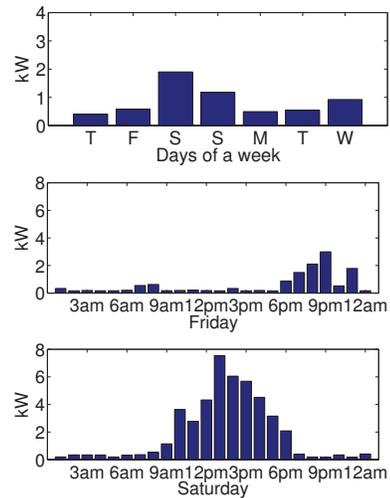


Figure 13: The daily and hourly electricity consumption of a household shows weekly and daily patterns.

Adopting the standard security practices is an effective solution. However, implementing such a defense requires the replacement of AMR meters with new meters or at least an upgrade of the firmware of all deployed AMR meters.

6.3 Jammer add-on

In systems deployed at the scale of million units, the cost of installation may outweigh the hardware cost of the devices. We are unaware of any remote firmware update capability for these meters. Thus, a firmware upgrade would require skilled maintenance staff to work on each meter. To substantially reduce the cost of such an upgrade, it is possible to package a protection mechanism into a separate add-on device, which can be physically attached and secured to a meter by lower-skilled personnel than a firmware upgrade.

The central component of this add-on device is a Privacy Preserving Jammer (*PPJ*), which can prevent continuous RF eavesdropping on packets in plaintext by masking meter transmissions. The PPJ continuously monitors channels and emits a jamming signal immediately after it detects a packet transmitted by the target meter to prevent eavesdropping. Meanwhile, to allow drive-by or walk-by meter reading, the PPJ can be temporarily deactivated remotely by authorized meter readers for a period just long enough to allow privileged meter reading.

Jamming Parameters. To reduce the complexity and cost, PPJ utilizes a narrowband transceiver that can listen or transmit only on one channel. The PPJ cycles through the meter's channel hopping sequence and emits a protocol-specific jamming signal to mask the data packets. Note that the proper channel sequence can be identified by searching for transmissions on each channel during initialization, or it can be acquired as prior knowledge from meter companies.

To effectively obscure AMR data packets, the PPJ transmits over the packet channel bandwidth (200 kHz) (shown in Fig. 14) for the entirety of a packet (5.8 ms). The jammer power should be larger than the meter's, but within FCC regulation. By mounting the PPJ close to the meter, the eavesdropper's ability to decode packets becomes inde-

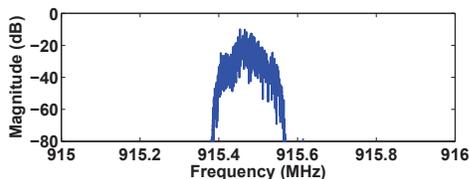


Figure 14: FFT of a PPJ jamming signal.

pendent of their location. Given any far-field eavesdropper location, the jammer-to-signal (J/S) ratio at the eavesdropper will always be greater than 1 (i.e., $J/S > 0$ dB), which represents ample distortion to prevent OOK communication (i.e., $BER = 0.5$) [29].

The PPJ minimizes its interference with other devices operating in the 915-MHz ISM band and obeys FCC regulations. The FCC limits the peak output power for frequency hopping to 30 dBm and places no restriction on duty cycle [30]. The PPJ transmits at a power level less than 0 dBm (yet still greater than the AMR meter’s transmission level). It overlays its jamming signals with meter transmissions and remains silent when no data is transmitted, thus greatly reducing any interference.

Deactivation Protocol. To support drive-by meter reading, the PPJ can be deactivated by an authorized reader, following the three-way handshake protocol illustrated in Fig. 15. The protocol requires three messages, and the third one ensures that the drive-by truck confirms the deactivation before the PPJ stops jamming. All three messages use basic signatures for authentication, and they are exchanged on the control channel (an unused channel around 915 MHz, e.g., 914.2 MHz). To ensure reception of the first deactivation request message, the PPJ can periodically switch to the control channel for listening, and the deactivation message can employ a preamble long enough for the PPJ to detect.

Experiments. We conducted defense experiments in the same apartment as the eavesdropping experiments in Section 4. We programmed one USRP as an RF sniffer and one as a PPJ. As before, an antenna for the RF sniffer was mounted against a window to collect data from the author’s meter. The PPJ was placed close to the author’s meter and approximately 3m away from the RF sniffer. We studied our defense strategies against two RF sniffers: one is a narrowband RF sniffer hopping through channels and the other is a wideband sniffer that monitors all channels simultaneously.

Using a narrowband RF sniffer, without turning on the PPJ, we were able to receive about 2 packets per hour on each channel, as shown in Fig. 16 (a). Once we turn on the PPJ, no packets can be received on any channels. To study the relationship between jamming bandwidth and the pph, we implemented a bandlimited jammer, which continuously jammed at 910-920 MHz, as shown in Figure 16. We observed that we blocked *all* packets in this band, and reduced the number of received packets in neighboring channels.

7. RELATED WORK

Non-intrusive Load Monitoring (NILM). NILM systems monitor the total load at an electric meter to extract individual appliance profiles. NILM algorithms can be divided into two categories based upon the signatures they use: steady-state and transient [26]. Transient techniques

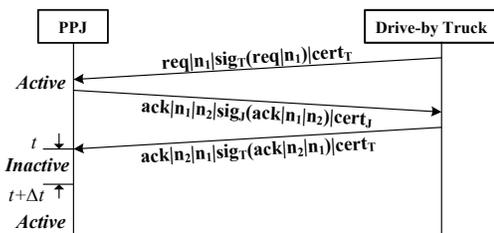


Figure 15: An example deactivation protocol for disabling the PPJ briefly for meter reading.

require high frequency measurements (e.g., Msp/s) [31, 32], while steady-state techniques utilize low frequency measurements and perform edge detection to identify appliances [26, 33]. Recent work examined power consumption in the frequency domain [34], extending the capabilities of traditional transient solutions by empowering differentiation between similar appliances. We used prior work [26] to evaluate the privacy breach of AMR meters.

Consumer Privacy of Load Monitoring. Researchers have investigated privacy leakage by employing NILM systems. Mikhail et al.[5] proposed a method to infer a resident’s activities from demand-response systems. They first employed an existing NILM algorithm to recognize the running time schedules of various appliances. Then, extraction routines were used to determine occupancy schedules, sleeping cycles, and other activities. In their earlier work[4], they investigated the impact of sampling rate on the accuracy of personal activity inference. They showed that even with 20-minute time resolution, attackers could still get meaningful estimates of a user’s activities with 70% accuracy. Our work complements theirs, since we proposed several practical attacks on deployed electric meters (both visual and RF-based) for acquiring consumption data, which can serve as data input to their study. To preserve consumer privacy from load monitoring, a protection system called NILL was proposed recently in [35]. They used an in-residence battery to mask the variance in load to counter potential invasions of privacy. We believe that our defense strategies against wireless attacks can complement NILL.

Reverse Engineering. Researchers have used reverse engineering methodology to expose security loopholes in systems when the designers tried to secure the system by obscurity. Rouf et al. [36] used a similar methodology to discover security and privacy risks of tire pressure monitoring systems. Nohl et al. [37] used reverse engineering to reveal ciphers from a cryptographic RFID tag that is not known to have a software or micro-code implementation. With some prior knowledge of the cipher, researchers used a black box approach [38, 39] for cryptanalysis of ciphers. Bortolozzo et al. [40] used reverse engineering to extract sensitive cryptographic keys from commercially available tamper resistant cryptographic security tokens by exploiting vulnerabilities in their RSA PKCS#11 based APIs.

Attack on Wireless Channel. Checkoway et al. [41] presented an analysis of vulnerabilities of automotive short-range wireless communications (Bluetooth), and long-range wireless communications (cellular). Francillon et al.[42] demonstrated relay attacks against keyless entry systems, and [43, 38, 44] also employed attacks on the RFID-based protocols used by engine immobilizers to identify the presence of an owner’s ignition key. Clark et al.[45] analyzed the security of

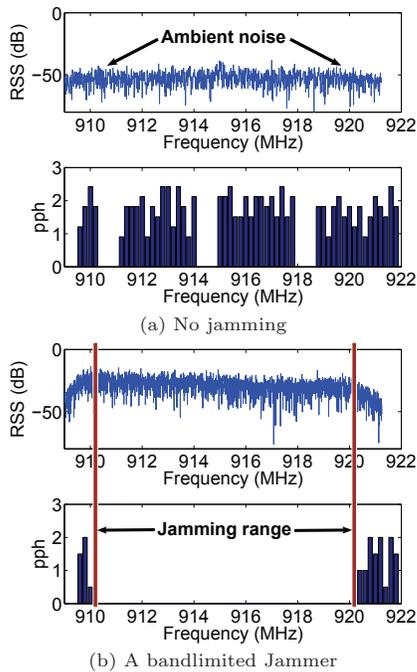


Figure 16: Channel spectrum plots (top) and received pph on each channel (bottom) for two scenarios.

P25 systems against both passive and active adversaries and showed that even when encryption is used, much of the basic meta-data is sent in the clear and is directly available to a passive eavesdropper. AMR systems being studied in this paper differ in several aspects from prior studied systems.

Defensive Jamming. Defensive jamming has also been proposed to protect medical devices [46]. Although sharing similar concepts, AMR meters involve a different physical layer technology (frequency hopping), which makes jamming harder to perform. Differing from prior work [46], the focus of our paper is to provide insight from both attack and defense sides.

8. CONCLUSION

AMR systems utilizing low power radio frequency (RF) communications have been widely deployed for automatically collecting utility consumption data. This work shows that currently deployed AMR systems are vulnerable to spoofing attacks and privacy breaches. Although AMR systems use frequency hopping, we were able to reverse engineer their communication protocol and launch spoofing attacks against a representative meter reader. Surprisingly, we found that AMR meters broadcast readings every 30s regardless of whether any ‘drive-by’ or ‘walk-by’ meter readers are in range, and meters have a communication range larger than expected. Through wireless monitoring, we harvested consumption data from 485 meters within a 300m radius region. This indicates that the millions of meters that have already been installed are at risk.

A few standard security remedies are available to cope with the discussed vulnerabilities. Yet, none of them are adopted in the deployed AMR meters that we studied. Adding those remedies to existing meters requires the upgrading of existing meters, which if too costly, can be replaced with an

alternative schemes that we call PPJ. It utilizes jamming to protect against the leakage of legacy devices and requires no modification of the deployed meters. Our pilot experiments offer a proof-of-concept that PPJ can be used to prevent information leakage of AMR meters.

Acknowledgement

The authors would like to thank anonymous reviewers for their valuable feedback and David Metts for assisting with the experiments. This work has been funded by NSF CNS-0845671.

9. REFERENCES

- [1] “Stop smart meters.” [Online]. Available: <http://stopsmartmeters.org/>
- [2] A. Bleicher, “Privacy on the smart grid,” *IEEE Spectrum*, October 2010.
- [3] N. Dallas-Fortworth, “Smart meters can be hacked: Security experts,” October 2009.
- [4] M. Lisovich and S. Wicker, “Privacy concerns in upcoming residential and commercial demand-response systems,” in *2008 Clemson University Power Systems Conference*. Clemson University, 2008.
- [5] M. A. Lisovich, D. K. Mulligan, and S. B. Wicker, “Inferring personal information from demand-response systems,” *IEEE Security and Privacy*, vol. 8, pp. 11–20, 2010.
- [6] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private memoirs of a smart meter,” in *2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys 2010)*, November 2010.
- [7] P. McDaniel and S. McLaughlin, “Security and privacy challenges in the smart grid,” *IEEE Security and Privacy*, no. 3, pp. 75–77, 2009.
- [8] D. L. Ulrich Greveler, Benjamin Justus, “Multimedia content identification through smart meter power usage profiles,” in *Computers, Privacy and Data Protection*, 2012.
- [9] P. M. S. Daniel J. Solove, Marc Rotenberg, *Privacy, Information, and Technology*. Aspen Publishers, Inc, 2006.
- [10] S. McLaughlin, D. Podkuiko, and P. McDaniel, “Energy theft in the advanced metering infrastructure,” in *Proceedings of the 4th international conference on Critical information infrastructures security*, ser. CRITIS’09. Springer-Verlag, 2010, pp. 176–187.
- [11] N. I. of Standards and Technology, “Guidelines for smart grid cyber security: Vol. 2, privacy and the smart grid,” *The Smart Grid Interoperability Panel’s Cyber Security Working Group*, vol. NISTIR 7628, 2010.
- [12] U. E. I. Administration, “Annual electric power industry report,” 2010.
- [13] Comverge, “Comverge integrates AMR metering technology into the converge IntelliSOURCE Platform,” 2010.
- [14] H. Ali, “Debunking the battery life expectancy myth between AMI and AMR,” September 2011. [Online]. Available:

- <http://www.waterworld.com/index/display/article-display/3002583591/articles/waterworld/water-utility-management/2011/09/debunking-the-battery-life-expectancy-myth-between-ami-and-amr.html>
- [15] A. Rial and G. Danezis, "Privacy-preserving smart metering," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*, ser. WPES '11, 2011, pp. 49–60.
 - [16] H. Khurana, M. Hadley, L. Ning, and D. Frincke, "Smart-grid security issues," *IEEE Security and Privacy*, vol. 8, no. 1, pp. 81–85, 2010.
 - [17] T. D. Tamarkin, "Automatic meter reading," *Public Power*, vol. 50, 1992.
 - [18] T. G. Paraskevakos, "Sensor monitoring device," United States Patent, 1972.
 - [19] N. E. T. Laboratory, "NETL modern grid strategy: Advanced metering infrastructure," February 2008. [Online]. Available: <http://www.netl.doe.gov/smartgrid/referenceshelf/>
 - [20] M. L. Grindahl and Q. S. Denzene, "Automatic/remote RF instrument monitoring system," January 1989.
 - [21] J. McNabb, "Vulnerabilities of wireless water meter networks," February 2011. [Online]. Available: https://media.blackhat.com/bh-us-11/McNabb/BH_US_11_McNabb_Wireless_Water_Meter_WP.pdf
 - [22] "Ettus Research LLC," <http://www.ettus.com/>.
 - [23] "Mini-circuits," <http://www.minicircuits.com>.
 - [24] J. Manweiler, N. Santhapuri, S. Sen, R. Roy Choudhury, S. Nelakuditi, and K. Munagala, "Order matters: transmission reordering in wireless networks," in *Proceedings of MobiCom '09*. ACM, 2009, pp. 61–72.
 - [25] A. Goldsmith, "Stanford University EE 359 Wireless Communications Course Notes," <http://www.stanford.edu/class/ee359/>.
 - [26] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.
 - [27] V. Brik, S. Banerjee, M. Gruteser, and S. Oh, "Wireless device identification with radiometric signatures," in *Proceedings of the 14th ACM international conference on Mobile computing and networking*, 2008, pp. 116–127.
 - [28] N. Patwari and S. K. Kasera, "Robust location distinction using temporal link signatures," in *Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, 2007, pp. 111–122.
 - [29] J. Proakis, *Digital Communications*. McGraw-Hill Science, 2000.
 - [30] F. C. Commission, "Operation within the bands 902-928 MHz, 2400-2483.5 MHz, and 5725-5850 MHz." *Section 15.247*.
 - [31] S. Leeb, S. Shaw, and J. Kirtley, J.L., "Transient event detection in spectral envelope estimates for nonintrusive load monitoring," *Power Delivery, IEEE Transactions on*, vol. 10, no. 3, pp. 1200–1210, 1995.
 - [32] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong, "Power signature analysis," *Power and Energy Magazine, IEEE*, vol. 1, no. 2, pp. 56–63, 2003.
 - [33] M. Marceau and R. Zmeureanu, "Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings," *Energy Conversion and Management*, vol. 41, no. 13, pp. 1389–1403, 2000.
 - [34] S. Gupta, M. S. Reynolds, and S. N. Patel, "ElectriSense: Single-point sensing using EMI for electrical event detection and classification in the home," in *Proceedings of Ubicomp '10*. ACM, 2010, pp. 139–148.
 - [35] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proceedings of CCS '11*. ACM, 2011, pp. 87–98.
 - [36] I. Rouf, R. D. Miller, H. A. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, and I. Seskar, "Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study," in *Proceedings of USENIX Security Symposium*, 2010, pp. 323–338.
 - [37] K. Nohl and D. Evans, "Reverse-engineering a cryptographic RFID tag," in *Proceedings of USENIX Security Symposium*, 2008.
 - [38] S. C. Bono, M. Green, A. Stubblefield, A. Juels, A. D. Rubin, and M. Szydlo, "Security analysis of a cryptographically-enabled RFID device," in *Proceedings of USENIX Security Symposium*, 2005.
 - [39] *From Fish to Colossus: How the German Lorenz Cipher was Broken at Bletchley Park*. Cragon Books, 2003.
 - [40] M. Bortolozzo, M. Centenaro, R. Focardi, and G. Steel, "Attacking and fixing PKCS 11 security tokens," in *Proceedings of CCS '10*, 2010.
 - [41] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage, "Comprehensive experimental analyses of automotive attack surfaces," in *Proceedings of USENIX Security Symposium*, 2011.
 - [42] A. Francillon, B. Danev, and S. Capkun, "Relay attacks on passive keyless entry and start systems in modern cars," in *Proceedings of NDSS Symposium*, 2011.
 - [43] T. Eisenbarth, T. Kasper, A. Moradi, C. Paar, M. Salmasizadeh, and M. M. Shalmani, "On the power of power analysis in the real world: A complete break of the KeeLoq code hopping scheme," in *Proceedings of N. Smart, Eurocrypt*, 2008.
 - [44] S. Indestege, N. Keller, O. Dunkelman, E. Biham, and B. Preneel, "A practical attack on KeeLoq," in *Proceedings of N. Smart, Eurocrypt*, 2008.
 - [45] S. Clark, T. Goodspeed, P. Metzger, Z. Wasserman, and M. Blaze, "Why (Special Agent) Johnny (Still) Can't Encrypt: A Security Analysis of the APCO Project 25 Two-Way Radio System," in *Proceedings of USENIX Security Symposium*, 2011.
 - [46] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, "They can hear your heartbeats: non-invasive security for implantable medical devices," in *Proceedings of the ACM SIGCOMM*, 2011, pp. 2–13.