The Hitchhiker's Guide to Successful Residential Sensing Deployments

Timothy W. Hnat hnat@virginia.edu
Vijay Srinivasan vs8h@virginia.edu
Jiakang Lu jklu@virginia.edu
Tamim I. Sookoor tis5m@virginia.edu
Raymond Dawson rcd4j@virginia.edu
John Stankovic stankovic@virginia.edu
Kamin Whitehouse whitehouse@virginia.edu

Department of Computer Science
University of Virginia
Charlottesville, CA 22902

Abstract

Homes are rich with information about people's energy consumption, medical health, and personal or family functions. In this paper, we present our experiences deploying large-scale residential sensing systems in over 20 homes. Deploying small-scale systems in homes can be deceptively easy, but in our deployments we encountered a *phase transition* in which deployment effort increases dramatically as residential deployments scale up in terms of 1) the number of nodes, 2) the length of time, and 3) the number of houses. In this paper, we distill our experiences down to a set of guidelines and design principles to help future deployments avoid the potential pitfalls of large-scale sensing in homes.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Distributed Networks; C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and Embedded Systems

General Terms

Design, Experimentation

Keywords

Deployment, Buildings, Smart Homes, Sensor Networks

1 Introduction

On average, Americans spend 65% of their time in their own homes [1]. Residential sensing systems can collect a wealth of information about people, much of which cannot be gathered in any other way. Sleep patterns, eating habits, hygiene, and many other characteristics of a person's lifestyle manifest primarily in the home, and statistics about these activities of daily living (ADLs) are valuable indicators of health. Homes also constitute a large fraction of the carbon and energy footprint over which a person has direct control, and statistics about electricity, water, and product consumption in the home are valuable for personal conservation efforts. Finally, homes are the core of personal and family life, and residential sensing can be valuable for assisting, coordinating, and tracking personal activities such as shopping, meal preparation, and vehicle sharing for after-school activities.

Goals: In this paper, we discuss the challenges of designing, deploying, and managing large-scale residential sensing systems. The goal is not to make a scientific contribution, but rather to summarize lessons learned from the experience of deploying over 1200 sensors in over 20 homes over the course of several years, including both commercial off the shelf (COTS) devices and our own custom designs. Each of these deployments was performed for a different application or experiment, and in this paper we review all deployments at once to distill out a single set of guidelines and design principles supported by data, pictures, and anecdotes from our own experience. We focus on common myths and misconceptions to help future deployments avoid potential pitfalls. Most of the guidelines that we identify are simple common sense, but are obvious only in hindsight; the contribution of this paper is the collection of hazards and challenges that make these guidelines appear obvious.

Findings: The key finding of our study is a *phase transition* in which deployment effort increases discontinuously as residential deployments scale up. Deploying a single sensor in a home is nearly effortless: plug it into a wall socket and wirelessly relay the data to a local computer or remote server. We pushed the scale of residential sensing deployments in three dimensions 1) the number of nodes 2) the duration of the deployment through time, and 3) the number of homes in which the system is deployed simultaneously. We found that each of these dimensions introduces its own challenges and that, at scale, indoor sensing systems require no less design and preparation than outdoor deployments. Some of the challenges that we encountered indoors are identical to those

SenSys'11, November 1–4, 2011, Seattle, WA, USA.

already observed in outdoor deployments [2]. Other challenges however, were entirely unexpected or were analogous but not identical to those outdoors. Our experience helped to elucidate these similarities and draw analogies where possible. As we scaled up our deployments, some of our more surprising findings included:

- Power became a scarce resource once the number of sensors exceeded the number of 120V wall sockets in each home (typically 20-30). Furthermore, wall-powered nodes were 2.3x more likely to lose power than battery-powered nodes.
- Wireless connectivity in homes was worse than expected because sensors were often placed in exceptional locations.
- Children, pets, and robotic vacuums became *"environmental hazards"* once the number of sensors grew to cover a substantial fraction of surfaces, appliances, and fixtures in the home.
- Houses became *"remote environments"* with limited physical access once the deployments scaled to homes other than the investigators' own.
- User participation dropped precipitously as the time duration of experiments grew.
- Aesthetic appeal became a limiting design factor once the number of sensors grew large enough to be visually salient, and time durations exceeded a few months.
- Maintenance time grew quickly as we used a greater number of different COTS platforms, because each new platform introduced a new set of possible failures.

2 Related Work

Researchers have been deploying sensors in homes for several decades, but few deployments reveal the same insights about residential deployment challenges that ours do. Most large-scale systems deployed to date have used a permanent installation of sensing infrastructure such as in-line power, communication lines, and secure mounting locations for the sensors. In contrast, our deployments focus on infrastructure-less systems that use surface-mounted sensors, wireless communication, and wall sockets or battery power. Infrastructure-less systems dramatically reduce installation cost and we therefore expect them to be the most common type of smart home in the future, particularly for do-it-yourself installations or for applications like home energy management (HEM) where the financial benefits do not justify costly infrastructure. In recent years, infrastructure-less sensing has become common, but deployments to date have been limited in number (a few dozen sensors) and/or duration (days to weeks). To our knowledge, our deployments are unique in that they scale up infrastructure-less residential sensing in three ways simultaneously: 1) hundreds of sensors per home 2) multiple months of deployment time, and 3) multiple homes.

Infrastructure-based deployments: Many residential sensing systems have been deployed with hundreds of sensors for long time durations, but the homes were specially equipped with data and communication infrastructure. For example, the Neural Network House [3] at the University of Colorado at Boulder is a long-term testbed in use since the mid-90's with an impressive array of instrumentation that senses and controls light intensity, sound levels, temperature, motion, the status of doors and windows, ceiling fans, space heaters, furnaces, and water heaters. To implement the system, a historical building was retrofitted with nearly five miles of low-voltage conductor to collect sensor data and a power line communication system to control actuators. The Aware-Home [4] is a smart home designed at the Georgia Institute of Technology during the late-90's. The duplex home was constructed from scratch and contains two identical living spaces and a control room for centralized computation. The building uses a high-bandwidth network to maximize the information provided to occupants or their caregivers, and uses an extensive sensor suite including cameras and microphones. MavHome [5] is a long-term testbed deployed at the University of Texas at Arlington since the early 2000's. It acts as an intelligent agent with the goal of maximizing occupant comfort while minimizing the operating cost of the home. MavHome uses a wide array of sensors and controllers that are supported by in-line power and communication over CAT5 cables. The PlaceLab is a laboratory designed at MIT in the mid-2000's for temporary living by study participants. It has hundreds of sensors built into the walls, fixtures, and cabinetry, including cameras and microphones [6]. "Practical limitations" to installing portable sensors in real homes were cited as a motivation for the construction of the lab, but a detailed description was not given.

Infrastructure-less deployments: Several ubiquitous computing studies have used infrastructure-less sensing for studies in real homes, but were limited to a few dozen sensors and/or a few days or weeks of deployment duration. For example, Tapia et al. deployed 77 contact sensors throughout a house for two weeks [7]. Cook et al. deployed 18 motion sensors and 2 temperature sensors for four months [8]. Kasteren, et al deployed 14 state-change sensors for four weeks [9]. Most commercial home automation systems are infrastructure-less. Early systems by X10 plugged into wall sockets and used power line communication, while later systems are battery operated and wireless. Newer and more expensive systems, such as those that support the Z-wave protocol, are also designed to operate without a computer in the home to act as a control station. These systems are installed in millions of homes, but typically have only a few dozen sensors. In Section 5, we describe several reasons why these systems are difficult to scale to hundreds of nodes.

Lessons learned: Many of the lessons discussed in this paper have undoubtedly been observed during other home sensing projects, but other studies did not result in a comprehensive listing of these lessons or a distillation of guidelines from which future projects could benefit. Papers that do reflect on guidelines do not focus on deployment and system operation, as this paper does. For example, Edwards and Grinter discussed seven challenges that must be overcome in order for smart homes to be viable [10], including possible social ramifications, the effect on home life, the challenges

	#Homes	Weeks	Motion	Object Use	Door Height	Wearable Tracking	Light Switch	Power (Plugs)	Power (Circuits)	Power (Mains)	Water Mains	Custom Thermostat	Active Register	Light/Temp Humidity
A	11	1-2	25-30	12-20	-	-	-	-	-	-	-	-	-	-
B	1	1	-	-	12	12	-	-	-	-	-	-	-	-
C	3	3-4	15-25	-	-	-	-	-	-	-	-	-	-	12-25
D	1	2	4	-	-	-	-	2	-	-	-	-	-	-
E	1	2	5	-	-	-	-	-	1	-	-	-	-	-
F	1	28	65	13	13	16	22	-	-	1	1	1	12	86
G	1	44	54	7	31	14	22	8	37	1	1	1	12	29
H	1	39	15	7	14	-	11	4	-	1	1	-	-	-
I	1	32	25	10	30	-	31	3	48	1	1	-	-	8
J	1	25	14	5	17	-	7	2	-	1	1	-	-	-

Table 1. Over the course of several years, we deployed 100's of sensors in over 20 homes, ranging from small, short-term to large-scale and long-term deployments. This table summarizes our main deployments. One home was used for deployments A, B and H, and another home was used for both deployments F and G.

of ambiguous data, and the lack of a system administrator. They also discuss challenges due to the piece-meal addition of technology to homes, the need for device interoperability, and the complexity of designing heterogeneous systems. In "Principles of Smart Home Control" [11] Davidoff et al. attempt to rephrase the traditional question of "How can smart home control systems help users regain control of their devices?" to "How can smart home control systems help families regain control of their lives?" The authors provide design principles that help focus design efforts on the targeted audience rather than the devices being controlled. In contrast, our paper focuses on practical guidelines intended to help achieve reliable system operation during large scale, long duration, infrastructure-less deployments.

A set of guidelines with a similar goal was summarized in a "Hitchhiker's guide for WSN deployments" [2]. However, that guide focused on the challenges of outdoor deployments, many of which are not relevant to indoor deployments. For example, indoor environments are not subject to extreme weather conditions or temperature fluctuations that can affect clock drift, battery lifetime, and wireless connectivity. Indoor sensors are typically one hop from a base station and so they perform little if any distributed processing. Therefore, it is much less important to have remote control over sensors and visibility into their internal computations, or to use simulation to test and debug distributed protocols. Finally, there is no clear distinction between a testbed and a deployment: a single home serves both purposes. On the other hand, there are also many similarities between indoor and outdoor deployments: homes can have limited physical access, similar to *"remote"* outdoor environments; LEDs must be turned off, but for occupant comfort rather than for energy efficiency; wild animals are not a threat indoors, but children, guests, and robotic vacuums can be hazardous; preparation and organization before deployment is necessary, but with a very different checklist (bring a hand vacuum rather than an ice-axe). In this paper, we reflect on our indoor deployments and identify similarities or draw analogies to outdoor deployments when possible in order to strengthen and generalize the lessons learned from both, even if they are not obviously comparable at first.

To our knowledge, this paper is the first compilation of potential pitfalls, common misconceptions, and practical advice about indoor sensing in residential environments.

3 Overview of Our Residential Deployments

Over the past several years, we have deployed over 1200 sensors in over 20 homes in a series of residential sensing experiments. The scale, duration, and the number of sensors in each deployment varied from home to home, year to year, and experiment to experiment, but the goal of all of these sensing systems was to monitor human activity in the home. Our deployments have collected over 17 billion data points portraying various aspects of over 25 people's lives, including the use of appliances, water fixtures, lights, energy, doors, and individual rooms. The data has been used for numerous scientific studies [12, 13, 14, 15], many of which are still ongoing, and as of this writing 5 deployments continue to operate and collect data. The studies themselves are out of scope for this paper. In this section, we describe the sensors and the system architecture that we deployed in order to provide more focused context and perspective for the analysis and design guidelines that follow.

Sensors and Controllers: The sensors and actuators in our system changed with each home, each experiment, and each generation of hardware. A summary of our main deployments is shown in Table 1, and images of some of the sensors are shown in Figure 1. The full sensor suite includes over 200 different enclosures in a single home, many of which contained several different types of sensors, for a total of well over 350 different sensors.

Almost all deployments had at least one wall-mounted motion sensor per room (Figure 1(d)). We used motion sensors manufactured by X10 because they were inexpensive ($5 per sensor) and could be surface mounted using double-sided tape. Depending on the visibility required for the experiment, motion sensors were placed on walls, on walls next to doorways, on both sides of every doorway, and/or on every window. In one deployment, multiple simultaneous experiments resulted in 65 motion sensors in a home with 9 rooms.

We used contact reed switches (latch sensors) to detect the use of objects, including appliances, cabinets, doors/windows, lights, and water fixtures (Figure 1(h). We originally used X10 sensors due to cost. In later experiments when the contact switches served as ground truth sensors, however, we switched to Aeon Labs Z-wave devices that use a reliable communication protocol. We found reed switches to be too obtrusive to use on light switches for long-term deployments, and so we iterated through a custom design based

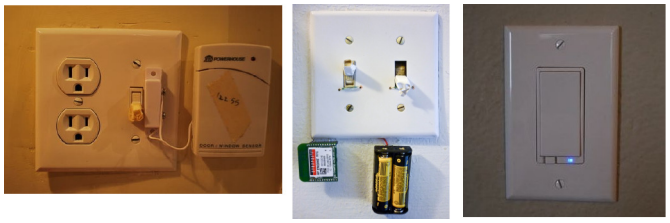

(a) Three Generations of Light Switch Sensors

(b) Plug Load Monitor

(c) Three Generations of Doorway Sensors

(d) Motion Sensor

(e) Three Generations of Active Registers

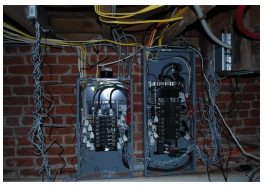

(f) Power Meter

(g) Water Flow Meter

(h) Reed Switches on Water Fixtures

(i) Light Sensor on Window

(j) Temperature Sensor

Figure 1. We deployed a wide array of sensors in homes, a subset of which are shown in this figure. We used commercial products whenever possible and designed and integrated custom solutions when necessary. We used latch sensors designed for doors and windows to detect use of other objects in the home, including appliances, cabinets, light switches, and water fixtures (h). As we gained experience and as the requirements of our experiments changed, we used multiple generations of hardware designs for several of the sensing sub-systems, including light switch sensors (a) doorway sensors that measure occupant height, motion, and door open/close status (c) and active registers that control air flow into each room (e).

on the synapse-wireless SNAP mote, and off-the-shelf GE Z-Wave home automation switches (Figure 1(a)).

We deployed ultrasonic range finders above doorways to measure the height of people as they walked throughout the homes to provide weak biometric identification, and designed three hardware generations using range finders manufactured by GoMotion, MaxBotix, and PING (Figure 1(c)). To evaluate our tracking system, we used wearable RF beacons and the motetrack tracking software [16].

For energy metering, we used the TED 5000 for whole-house power metering, the Powerhouse Dynamic eMonitor for circuit-level monitoring (Figure 1(f)), and an Aeon labs Z-wave meter for plug load monitoring (Figure 1(b)). We used Shenitech's ST301 transit-time ultrasonic flow meter for water metering and hot water usage monitoring (Figure 1(g)). We used different platforms to monitor light, temperature, and humidity levels, depending on requirements for sensor accuracy, sampling frequency, and battery lifetime, including telosB motes, La Crosse Weather Direct TX60U-IT weather sensors, and Onset data loggers.

For actuation, we used a Web-enabled thermostat from the BAYweb company to control the heating and cooling equipment. We designed and deployed three generations of *active air vent registers* that could be wirelessly opened and closed to control air flow into each room individually. We also designed and built a custom thermostat based on the Synapse-wireless SNAP device that used relay circuits to control the HVAC equipment and in-line duct dampers.

System Architecture: Our system consisted of over a dozen *sub-systems:* groups of sensors that interoperate and rely on the same power source, software stack, and wireless bridge or communication path. Most sub-systems were made by different manufacturers, e.g. X10 sensors, although some were custom designs. Figure 2 illustrates the integration of several sub-systems, including their power and communication resources. Because of this sub-system heterogeneity, each failure mode in the system caused data loss in a different subset of sensors.

Some of our sensing sub-systems were robust to both short-term power outages and broadband disconnections. Our telosb-based light and weather systems were completely battery powered and transmit data using a wireless bridge connected by USB to the gateway machine, which was typically a laptop computer with up to 5 hours of batter life and over 100GB of hard disk space for buffering data before it was sent to a remote database server hosted in our lab.

Many of our sensing sub-systems did rely on AC power at some point, and therefore lost data even during short-term power outages. For example, our active registers achieved very low power wireless communication by exploiting a communication backbone that was plugged into wall sockets. The X10 sensors and Weather direct weather sensors were all battery powered but used AC power for the wireless bridge. Our 3rd generation light switch sensors used in-line power from the lighting circuits. The doorway sensors, motetrack beacons, water meters, power meters, and thermostats were powered through wall sockets, sometimes using low-cost 120VAC to USB (5VDC) converters.

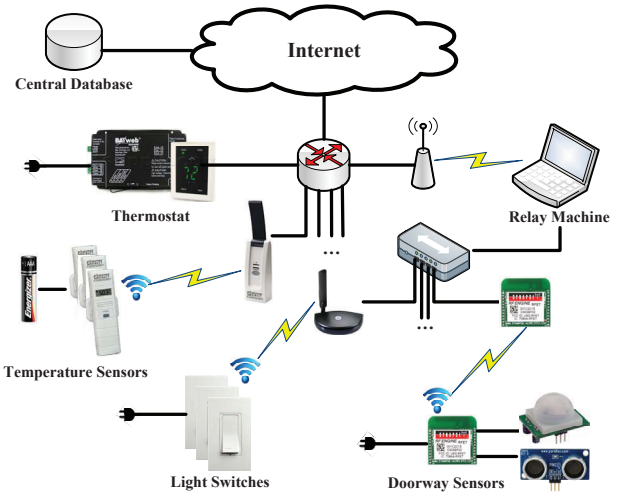

Figure 2. Our system architecture includes a heterogeneous array of sensing sub-systems, each of which is subject to a different set of failure modes.

Two of our sensing sub-systems lost data if the home's router or broadband connection failed because they connected directly to the home's router and sent data to a server hosted by the vendor, from where it was retrieved by the home gateway machine. These sub-systems include the Web-enabled thermostat and the Weather Direct weather sensors. Any sub-system failed if its own bridge or software stack failed, and all sub-systems failed if the gateway failed.

4 Reliability and Failure Analysis

In this section, we examine the main failure modes of our system during four recent deployments.

Failure Detection and Classification: We identify down time intervals for each sensor by defining the longest acceptable time interval τ between two consecutive data points; any interval longer than τ with no data is considered a down time interval for that sensor. Due to timestamp jitter, this parameter is set to be about five times larger than the sampling period for all sensors that collect data periodically. For motion sensors and object use sensors that are event driven and generate data only in response to occupant activity, this parameter is set to 36 hours. Table 2 summarizes the parameter τ used for each sensor type.

Once all down time intervals are identified for each sensor, we identify the root cause of failure. To do this, we exploit the fact that each failure mode of our system causes data loss in a distinct subset of sensors, identifying the root cause of each failure based on the set of simultaneous sensor failures, as follows:

1. **Wireless link loss:** down time of a single wireless sensor for less than 4τ[1]
2. **Battery dead:** down time of a single battery-powered sensor for longer than 4τ

[1]Wireless link loss is not assessed for event-driven X10 and Z-wave sensors because they do not transmit periodically.

Sensor Type	τ (in seconds)
Bayweb	3600
Water	2
Weather Direct	300
SNAP	60
TED	4
Light	120
E-Monitor	10
X10	129600
Z-Wave	129600

Table 2. We identified periods of down time for each sensor type by defining the longest acceptable time period τ between two consecutive data points.

3. **Plug disconnected:** down time of a single plug-powered sensor for longer than 4τ
4. **Sub-system down:** simultaneous down time of all sensors in a single sensor sub-system
5. **Internet Down:** simultaneous down time of all sensors reliant on a broadband link
6. **Power outage:** simultaneous down time of all sensors reliant on AC power
7. **Gateway down:** simultaneous down time of all sensors

If a down time interval satisfies more than one rule, only the root cause that explains the largest number of simultaneous sensor failures is asserted. For example, if all plug-powered sensors are down, each individual sensor failure could be explained by either a plug disconnection and a power outage. In this case, the system only asserts the power outage failure because it explains a larger number of sensor failures. This policy imposes a partial ordering on the failure explanations, and that ordering can be derived for each house based on the sensors installed in the house (Table 1) and the power, communication, and gateway or Internet resources used by each sensor (Section 3).

Our approach to identifying the root cause of failures is similar to that of Sympathy [17]. The key difference is that our approach is designed to run *post-facto* using only data loss to identify failures; it does not rely on meta-data collection about system operation. During system operation, we did use custom scripts and a tool called Nagios to identify failures and report them to the researchers, serving a purpose similar to Sympathy. Due to its on-line nature, however, this system suffered from false failure detections or mis-classifications due to short-term data delays. The post-facto approach that we present here uses hindsight to improve failure analysis.

Results: We executed the failure detection and classification algorithm described above on the four deployments named *G*, *H*, *I*, and *J* in Table 1 for the seven-month period from January 1, 2011 to August 1, 2011. We measure the effect of a system failure in terms of sensor down time, using a metric we call *sensor-days*: the length of the failure in days, multiplied by the number of sensors affected by the failure. This metric measures the importance of a system failure in terms of the number of sensors that are taken down, and gives all sensors the same level of importance. This metric avoids giving a higher weight to sensors that sample data periodically than to sensors that are event driven.

Root Cause	House G	House H	House I	House J
Sensing Sub-system	4107	642	4757	274
Gateway Down	5596	0	3	136
Plug Disconnected	509	30	474	10
Battery Dead	452	17	168	0
Wireless Link Loss	410	0	122	1
Internet Down	251	97	178	9
Power Outage	21	0	87	2

Table 3. The total sensor down time for four of the deployments listed in Table 1, broken down by root cause and measured in *sensor-days*: #days * #sensors.

Table 3 shows the sensor down time due to each type of system failure in each house, and Figure 3 shows how a subset of these failures change over the seven month period. This data illustrates several surprising trends. For example, AC power plug disconnections account for 1018 sensor-days of sensor down time while dead batteries account for 636 sensor-days. These 4 homes, however, contained 135 battery-powered sensors and only 93 sensors plugged into wall sockets. Thus, in our deployments, sensors were 2.3x more likely to lose data due to being unplugged than to battery failure. This statistic does not include any additional data lost by wall-powered nodes due to power outages.

Our analysis indicates that power outages were a minor cause of data loss, but this result can be misleading because most data lost during a power outage was attributed to gateway failure once its battery power was exhausted. Houses G and H both had frequent power outages, sometimes for over 24 hours, but the gateways in these houses only had battery power for 60 and 30 minutes, respectively. House I had much longer battery life on the gateway, causing more data loss to be attributed to the power outages even though it had far fewer power outages that the other two houses.

Wireless link loss is often a main challenge for reliable data collection, but in our deployments accounts for only 532 sensor-days of down time. This value is lower than other sources of sensor down time. For example, a single failure of the gateway's hard drive in House *G* caused the machine to be taken down, diagnosed, and reconfigured twice, causing 5590 sensor-days of sensor down time.

Failure of entire sensing sub-systems caused the largest amount of down time in all homes, accounting for 9780 sensor-days of down time across the four homes. These failures had many causes. For example, failure of the Z-wave or X10 wireless bridge would cause down time in the entire Z-wave sub-system. Similarly, a software bug would crash the software stack on the gateway that read data from the wireless bridge. Configuration changes were also common when the gateway, USB hub, or the home router would restart, for example after a power outage. In this case, the software stack might fail to restart, or the wireless bridge might acquire an new USB address and not be found by the software stack. Sub-system failures was so common because our system used nearly a dozen different COTS subsystems, each of which failed infrequently but in aggregate caused substantial down time. Sub-system failures became

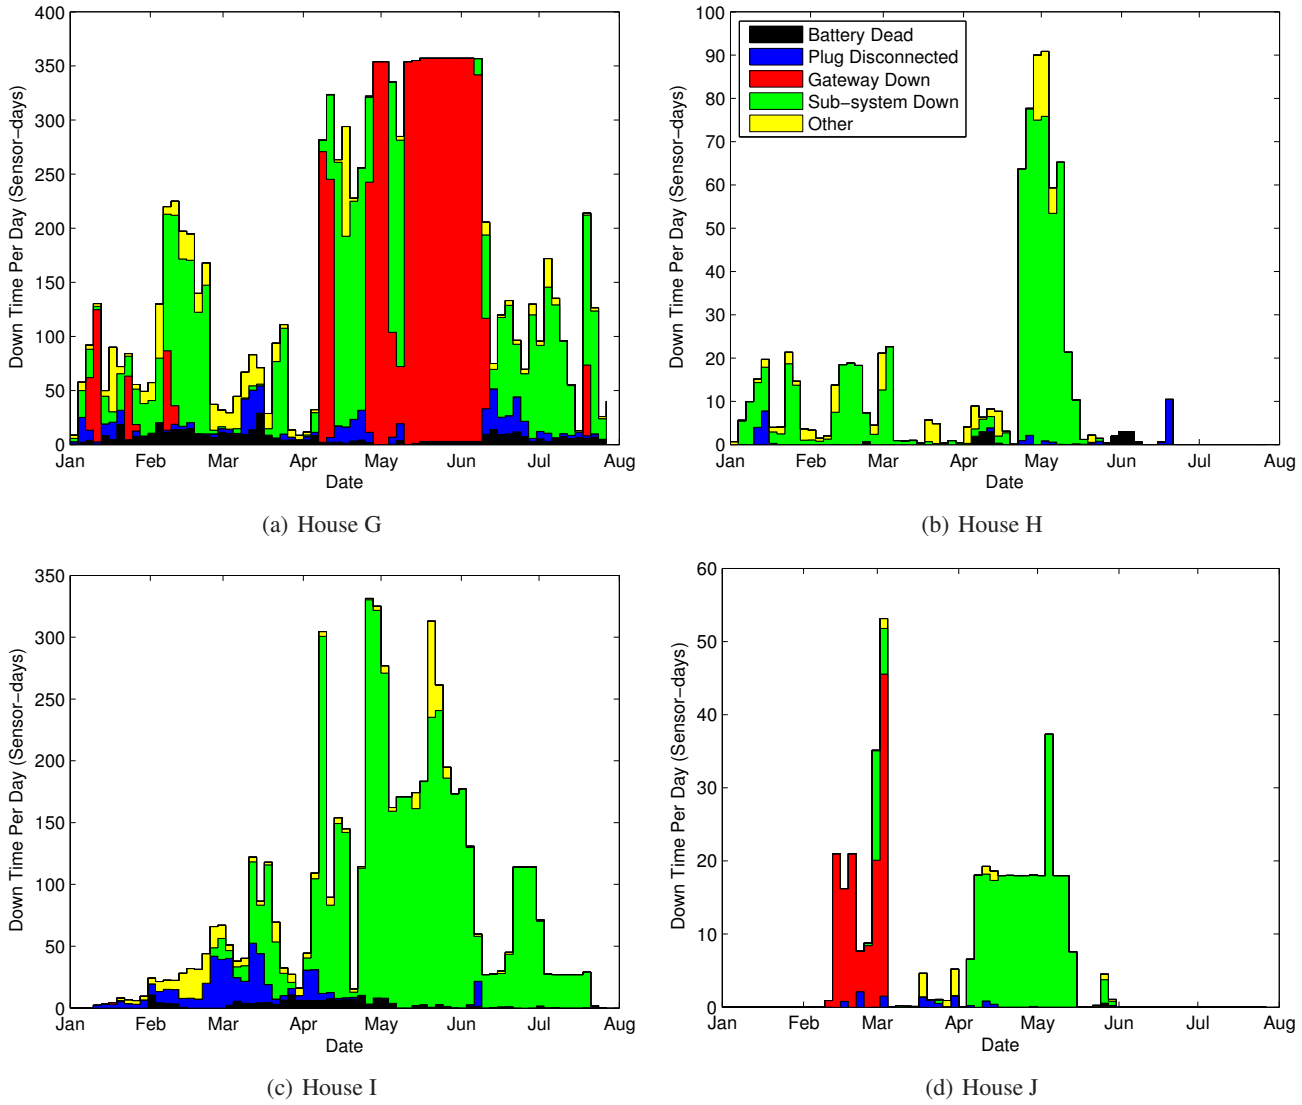

(a) House G

(b) House H

(c) House I

(d) House J

Figure 3. Total sensor down time for four deployments, broken down by root cause over time. Individual sensing sub-system failure is the dominant cause of down time, followed by failure of the gateway machine.

worse in all homes after April 8th, the submission deadline for this conference paper, when other priorities ensued and the research team did not maintain these sub-systems with the same urgency. This trend underscores the high frequency of maintenance required to operate the system.

5 A Hitchhiker's Guide

Based on our experience and analysis, we have created a set of guidelines to avoid many of the pitfalls and failures that we observed in our deployments. We believe that studies will increasingly be deployed in multiple houses over long time periods because every home and every person is different and, even within the same home, patterns change dramatically over the course of weeks, months, and even years. Furthermore, these systems will also include an increasing number of sensors in each home for redundant sensing because data validation becomes increasingly difficult over long time durations as users have a lower tolerance for surveys and continuous self-reporting. For these reasons, our guidelines focus on those challenges that are exacerbated at scale as deployments grow in terms of the number of sensors, the number of homes, and time duration.

5.1 Homes are Not a Power Panacea

Myth: *Sensors in homes can easily be powered using the wall sockets.*
Fact: *Wall sockets provide neither abundant nor reliable power, especially when deploying hundreds of nodes.*

The availability of 120VAC power in homes does simplify some sensing tasks, but it does not eliminate power issues and considerations. We explored three different ways of powering sensors in homes and found that batteries still have an important place in large-scale residential sensing systems.

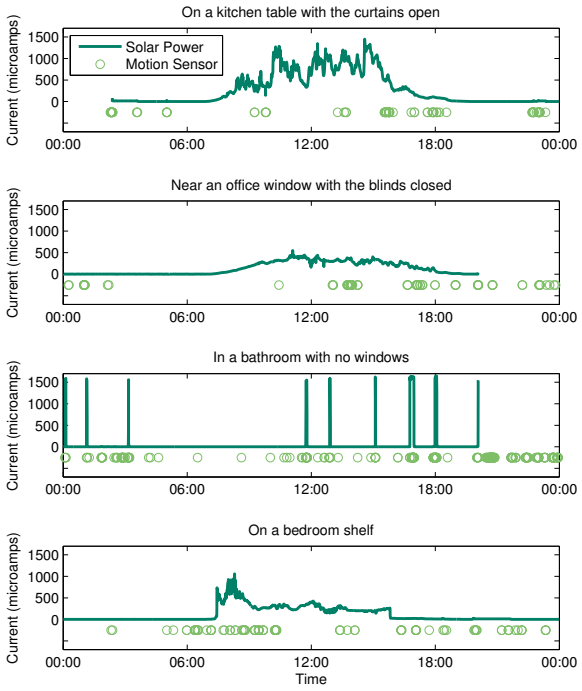

Figure 4. Solar energy can be harvested indoors for only 8-9 hours per day, despite much longer days and indoor light usage, and the energy harvesting time period does not necessarily correspond with times of human activity. The average power that can be harvested changes dramatically in different rooms and even locations within a room.

Wall Sockets: In our experience, several practical considerations severely limited the benefits of wall power, particularly when deploying at large scale. First, homes do not have enough wall receptacles to support all sensors in a large deployment, and when too many receptacles are used we found that users unplug sensors for their laptops, hairdryers, or vacuum cleaners. This became salient in one deployment when we used all but 6 out of 40 power receptacles in a house, many of which were filled beyond normal capacity using expansion adapters and power strips (Figure 5(a)).

Furthermore, contrary to the popular belief that wired power is more reliable than battery power, in our deployments a wall-powered node was 2.3x more likely to lose power than a battery-powered node. Wires were snagged, pulled, cut and disconnected. They are also of particular interest to animals and small children, and in one house many plugs were periodically pulled out by a robotic vacuum cleaner. Using wall power for a large deployment requires a vast amount of wiring, especially in older homes where receptacles are few and far between, and in one home we deployed over 250 linear feet of wire to power only 13 doorway sensors, 40 feet of which were for a single sensor, snaking through a hallway and into another room to reach the nearest receptacle. This wiring is both unsightly and a snagging hazard, particularly when it must run through doorways (Figure 5(c)) or near frequently moving objects (Figure 5(b)). It is inevitable to have periodic problems with this quantity of wiring, especially because it is particularly challenging to mount securely: in addition to the time and material cost for repair, long lines of holes from wire staples are unattractive and can be impossible to repair perfectly.

Not only do wall-powered sensors incur more data loss than battery powered sensors, they also require more frequent service calls. Over a 3 month period in 2 houses, we performed approximately two dozen service calls for 40 wall-powered sensors, but only 3-4 for over 100 battery-powered sensors. This is because battery powered sensors fail in a very predictable and correlated fashion, and all batteries can be changed at once after the first sensor fails. Thus, a single service call maintains all battery powered sensors. In contrast, wall-powered sensors fail independently, and each failure requires a service call. As the number of sensors increases, therefore, service calls for wall-power sensors become more frequent.

In-line Power: Sensors can be wired directly into the house wiring, avoiding the possibility of a plug disconnection. Snaking wires and opening up walls for sensor installation is expensive and undesirable unless doing a permanent and long-term installation, but some sensors can cheaply access in-line power if they are integrated into appliances or switch/receptacle enclosures. In-line power, however, can make software errors more difficult to recover from because cutting power to reboot the devices can only be done at the breaker box, in some sense *rebooting the house*. This is the main reason we opted to use batteries for our second generation light switch sensor (Figure 1(a)): as an experimental prototype, we expected the software to occasionally crash. Alternatively, a physical power switch or watchdog timer must be used.

In-line power is fairly reliable but is still subject to power outages, which can be frequent in any house due to blown fuses and electrical storms, and particularly in older houses with older wiring and above-ground power lines. In the 7 months that we analyzed, two of our houses (in the same neighborhood) both experienced about 2 power outages per month on average. Power outages did not typically last long, and accounted for a total of only 92 sensor-days of down time across all 4 homes. However, these outages often caused software and configuration errors that resulted in much longer down times. In comparison, most of our battery powered devices operate for a year or more without losing power. Thus, in-line power may incur fewer maintenance calls than battery power, but it is not necessarily more reliable.

Indoor Solar: Indoor solar power is one alternative to wall power and battery power. Figure 5.1 shows the waveform indicating the closed-circuit current produced by a 6 x 9 cm solar panel located at four different locations indoors, as measured by a Fluke 287 multimeter. This panel is a reasonable size for a wall-mounted sensor, if not too large.

Sensor	Average Power	Peak Power
Temperature	4.4 uW	0.6 mW
Motion	3.1 uW	8.1 mW
Object use	3.1 uW	7 mW
Height	18.1 mW	0.4 W
Active Register	1.8 mW	1.5 W

Table 4. The devices used in our deployments have a wide range of average and peak power characteristics.

The waveforms show that the amount of energy that can be harvested indoors varies greatly with the location of the solar panel. Thus, just as 120VAC is constrained to sensors that can be placed near AC wall sockets or integrated into a switch/receptacle, the use of indoor solar is constrained by the physical placement requirements of the sensors. Furthermore, the time of day that energy can be harvested does not typically correspond with the times at which the room is occupied, as indicated by the motion sensor readings on each graph, and the use of artificial lighting at night generates almost no current in any room besides the bathroom. This mismatch necessitates energy storage management, which complicates the use of indoor solar power.

The power generated by a solar panel depends on the actual load characteristics, but can be upper bounded by multiplying the closed-circuit current I_c by the maximum open-circuit voltage V_o, which was measured to be 0.8V. Thus, average power is upper bound by about 0.1mW, which is enough to power temperature/humidity sensors or motion sensors, but not height sensors (Table 4). For comparison, this same solar panel produced at most 102mW outdoors.

5.2 Homes Have Poor Connectivity

Myth: *Communication in homes can be achieved with single-hop wireless and/or power line modems.*
Fact: *Homes are small but can still be challenging RF environments, particularly for large-scale, dense, and heterogeneous networks.*

Homes are small geographically and can easily be covered by wireless technology such as WiFi. However, sensor nodes are often placed in exceptional locations, resulting in worse connectivity for sensors than is typical for laptops and WiFi. In one home, we placed thirteen 802.15.4 devices into metal duct work and 22 Z-wave light switches into metal junction boxes. Wireless connectivity of both radio protocols was dramatically reduced and many of these nodes had connectivity with only 1 or 2 neighboring nodes. Mounting nodes onto surfaces also introduces attenuation due to plaster, masonry, or concrete construction, and heavy metal appliances. We deployed in one house with concrete slab flooring separating the three levels and copper siding that produced a Faraday cage, isolating wireless sensors outside. Our current deployments produce up to 150 KB of data per second, which forces us to partition the network into multiple wireless channels, further inhibiting wireless connectivity. The network was further partitioned to separate the powered sensors from the unpowered sensors because they use different MAC protocols, and to separate devices made by different manufacturers. Throughout our deployments, we used five types of wireless networking schemes: one best-effort, single-hop wireless network; one proprietary single-hop network; one single-hop network with link-level reliability; and two networks with end-to-end multi-hop reliability. While the systems achieved fair delivery rates, link loss accounted for 532 sensor-days of down time over the seven month period, which is larger than expected over such short distances.

Power Line Communication: The electrical wiring in a home can be used as a transmission medium for high-frequency signals, providing another option for residential deployments called power line communication. However, power line communication is not always practical for low-power sensors. As discussed above, a large fraction of sensors in homes cannot connect to the home's electrical system due to the challenges of running new wires in the home. Furthermore, power lines are notoriously noisy, so narrow-band modems must make a trade off between cost, data rate, and robustness to noise. Robust power line modems that achieve 200 Mbps are commercially available, but cost more than many low-power sensors themselves. The power line modems typically used for low-power sensing devices such as X10 devices and the energy detective (TED) power meter are more likely to have low data rates and/or to be more vulnerable to noise on the line. This noise is caused by many electronic devices and can be difficult to avoid or filter because new devices are continuously added and moved throughout the house. On-line forums for X10 devices are rife with discussion of data loss as deployments scale to dozens of devices, and in our deployments we have observed up to two weeks of almost continuous data loss on the power line communication system used by our TED power meters.

5.3 Homes are Hazardous Environments

Myth: *Robust enclosures are only important for extreme outdoor environments.*
Fact: *Homes are safe environments for humans but can be hazardous for sensors, particularly when hundreds of sensors are deployed over long time durations.*

The calculus of mean time to failure (MTTF) applies to residential sensors just as it does to outdoor sensors: a low failure rate for a single sensor can translate to a high failure rate for dozens or hundreds of sensors. Through our experience, we have identified a number of unlikely causes of sensor failure in homes that, in the aggregate for hundreds or thousands of simultaneously deployed nodes, lead to weekly or even daily system maintenance. For example, sensors should always be child-proofed because they are both a curiosity and a choking hazard to toddlers and pets. Bright and flashing LEDs only exacerbate this hazard (Figure 5(a)). Sensors installed on objects such as microwaves, faucets, or light switches must be well secured to avoid being dislodged during normal use. For example, the wires on our first-generation light switch sensors were easily snagged (Figure 1(a)), and our faucet sensors were subject to both wet surfaces and user interference (Figure 1(h)). The mounting

(a) Overloaded Sockets (b) Snag Hazard (c) Snag Hazard

Figure 5. Wires are hazardous for reliability and maintenance. Users unplug devices when sockets are overloaded (a), and wires are often snagged, especially when they are near moving objects (b) or run through doorways (c).

techniques should not rely on users learning to accommodate sensors, because a large fraction of dislodged sensors in our deployments were due to guests, cleaning services, and other non-residents. Sensors and wires near the ground must be secured from brooms and vacuum cleaners, particularly of the robotic variety (Figure 5(b)). Sensors installed on furniture such as bookshelves may be moved or hidden and produce changed, invalid, or unreliable data.

Verify Constantly: With high failure rates and data rates, it is critical to quickly and automatically identify sensor failures. For example, during the course of several months we had over 500 sensors deployed in a half dozen houses, streaming on the order of 100 million data points per day. A 1-year mean time to failure per sensor translates to more than one failure per day and, indeed, only one day went by during that period with no sensor failures. To address this, we deployed a set of automated scripts for both component-level checks and end-to-end data verification. These include:

- *Network down:* ensure that the machine at each house can be reached.
- *Service down:* ensure that the service collecting a certain type of data at a house is running.
- *Last entry time:* ensure that each sensor has reported at least once within a certain period.
- *Minimum frequency:* ensure that each sensor is reporting with at least a minimum frequency.
- *Calibration:* ensure that the average value of a sensor is correct.
- *Time incorrect:* ensure that the local time on the machine at each house is correct.
- *Load high:* ensure that the CPU load of the server or machine at a house is sufficiently low.
- *Space low:* ensure that the disk space of the server or machine at a house is sufficiently high.
- *Timestamps incorrect:* ensure that the timestamps associated with a sensor's output is as expected.

These scripts were executed by a tool called Nagios that logged the results and reported them to the researchers. Getting researchers to respond to such alerts is still an open challenge, especially when there are several failures per day: it is important not to have too many or too few alerts. Our first implementation used email alerts, but a large number of false positives and transient failures caused researchers to ignore these alerts. We then used repeated emails every 10 minutes, but most such emails were spam filtered. Finally, we projected all critical alerts onto a wall in our lab, together with the duration of the alert (Figure 5.4). This approach was effective because the entire research team was aware of all failures, without the need for intrusive alerts. The projection was no longer actively used by the research team after April 8, which resulted in the large increase of un-repaired sub-system failures shown in Figure 3.

5.4 Homes are Remote Environments

Myth: *Maintenance visits are not a problem for homes.*
Fact: *Investigators have very limited access to deployments not in their own homes.*

Volunteers for sensor deployments must make a personal time commitment for scientists to enter their home and deploy or maintain the system. This time is precious and must be used wisely by the scientists. Short visits constrain the volunteer's mobility and appointment scheduling, and visits longer than 4 hours interfere with meals. Visits of a full day or more will typically be extremely limited. Therefore, deployment and maintenance visits must be highly efficient and optimized: dozens or hundreds of sensors must be deployed in a matter of hours. Configuring, installing batteries, assembling parts, and mounting a sensor may only take 10 minutes, but for 200 sensors this adds up to over 4 days of installation time, eight hours per day. Furthermore, sensor repairs must be made in batches to minimize service visits, which will increase the average time to repair. Sensor failures in researchers' homes were typically repaired in a matter of days while failures in volunteers' homes sometimes waited several weeks due to coordination constraints.

Custom Tailor Each Deployment: In contrast to previous outdoor deployments that used assembly lines and batch operations to minimize total deployment time [18], we found it also important to minimize *on-site* deployment time: system building, configuration, and testing must be moved to the lab, to the extent possible. Every house is slightly different, so each deployment requires two on-site components: the site visit and the deployment. During the site visit, scout out the deployment position of every single sensor, take measurements, photographs, and make records on a floor plan. These measurements must then be used for lab assembly and configuration of the sensors. If wires are to be run, the path of the wires should be determined and exact measurements made so that the wires can be cut to the length before soldering. In our case, we needed to check the location of pipes, the number and locations of electrical panels, styles of faucet fixture, the number of light switches and the number of switches in each gang box, the heights and widths of doorways, if the floors were wooden or carpeted, and if walls were plaster, concrete, or wallpapered. If accessing a crawl space, be sure somebody on site will be able to fit into it. Assess each location for wireless or wired connectivity options.

After the site visit, fully assemble and configure all sensors in the lab: give each node an ID and pre-determine its location. Print the locations on multiple floor plans so that multiple people can deploy the sensors in parallel. Put any labels with location and/or sensor ID on the back of the sensors so they are not eye-catchers for the users after deployment (Figure 1(d)). Cover LEDs before deployment; unlike programmable LEDs on experimental platforms, LEDs on COTS devices (Figures 5(a)) must be disabled or covered with electrical tape. For COTS sensors, insert batteries and remove any packaging material in the lab. Assign all deployment tasks before arriving on site. When designing and deploying, consider maintenance time requirements. For example, it may be faster to assemble sensors by hand, but debugging hand-soldered circuits on-site is challenging while taking them back to the lab requires two visits. Have circuits manufactured when possible, despite increased development time. Use velcro to attach sensors when battery compartments are only accessible from the back.

Any tools or items forgotten can add hours of delay during deployment. Put each sensor type in an individual box together with the deployment chart. If removing anything from the house, such as old light switches, bring extra boxes for simultaneous removal and installation. Bring jars for screws, wire nuts, and other small parts. Count the electrical sockets and bring expansion adapters or extension cables as necessary. Count the tools needed for each installation and bring enough sets for parallel installation. Bring tool belts to avoid putting tools onto fine surfaces. Bring flash lights, garbage bags, and a handheld vacuum cleaner for clean up afterward. Bring extra sensors to the site; some are sure to be broken.

Test Three Times: Several outdoor deployment studies have emphasized the need for both lab testing and deployment time validation [2]. In our deployments, we also found it critical to test immediately *after* deployment, because sensors that are designed to monitor people are greatly affected by the experimenters themselves. Thus, we advocate testing not just twice, but three times: before leaving the lab, at deployment time, and immediately after leaving the site.

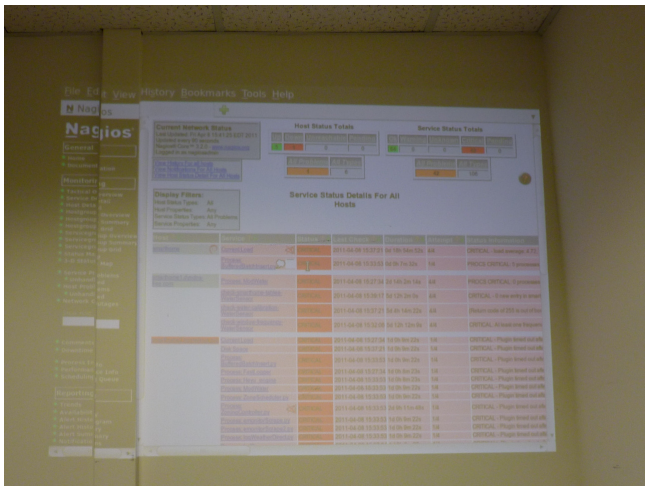

Figure 6. Our system constantly verifies system operation and critical alerts are project onto a wall in our lab.

In one of our deployments we installed over 60 motion sensors in a house and all motion sensors were transmitting constantly due to the number of researchers in the house, causing data loss due to wireless collisions and corruption. Once the deployment was done, however, the motion sensors responded normally. In another example, we installed our second generation doorway sensors into a house and confirmed that the system was working properly: all sensors responded with accurate height measurements and very little noise. However, immediately after the deployment far fewer people were in the house and noise levels skyrocketed due to an increase in ultrasonic multipath echoes.

5.5 Expect Limited User Participation

Myth: *Users can help maintain the system, and can provide validation data through surveys or questionnaires.*
Fact: *A user's ability to monitor and report activities in the home is limited by the need to do those activities, particularly in long-duration deployments.*

When scaling residential systems to a large number of houses and long-term deployments, it becomes increasingly difficult to collect *ground truth*: a true report of what really happened in the house, needed to validate results. Who cooked dinner on March 12, and at what time? When was the dishwasher turned on? Were the occupants sleeping, or just reading in the bedroom? Previous studies have used annotated video of people in a home [4], a human observer in the home [7], or controlled experiments with pre-determined activities [19]. However, these approaches do not scale well because they are very labor intensive and/or they interfere with the true patterns and characteristics of the natural residential environment.

An alternative approach is to ask the users to provide ground truth data through surveys or questionnaires. However, we have performed 5 such studies, each requiring a different level of user participation, and found that we could collect data either with high accuracy or for long time periods, but not both simultaneously. In real-time tracking studies that required constant participation, we could achieve high quality data for at most a few hours at a time. In studies that used surveys and self-reporting, users would report activity times with 1-minute precision several times per day for a few days, or with 15-minute precision once per day for a few weeks. Users could repair or report sensor failures for over a year, but would sometimes wait days or weeks before doing so. Interestingly, it was not feasible to infrequently query users about things like sequences of rooms occupied or light switch usage; even though the queries were very infrequent, it was too demanding to require people to continuously observe and remember such details about their own lives.

Much like the energy in a sensor's battery, user participation appears to come in finite quantities and can either be used intensely in short bursts or slowly over long periods. We do not believe that this trend is due to a lack of motivation. Would tracking, for example, have been more effective if we replaced the RF beacons with cell phones, which people are more motivated to carry? Our participants reported carrying the sensors as much or more often than their cell phones, and typically forgot to carry the sensors at the same times that they might not carry their phones: immediately after waking up, changing clothes, showering, or returning home. The length of time we could collect consistent data was not limited by the individual, but by aggregate group performance: in a multi-person home, it was highly likely that at least one person had not carried the device in any given time period. We found that limits on participation applied to co-investigators and non-investigators alike: it was not a lack of motivation, but rather that personal and family activities are a necessary part of life, even for co-investigators, and people can only tolerate so much interference due to participation in a residential sensing study.

Thus, ground truth validation is a potential pitfall of long-term, large-scale residential sensing studies, and may not be detected until weeks or months into a deployment. To address this, be sure to use redundant sensing and multiple ground truth techniques that can be validated against each other. This can lead to an explosion in the number of sensors. For example, one of our studies required only 2 sensors at the electrical and water mains of the house, but we needed to install over 100 sensors to validate our measurements by monitoring all light switches, plug loads, faucets and water fixtures, and major appliances. We also use these object sensors as a proxy metric for other studies such as tracking accuracy: the consistency between object use in the house and a person's predicted location. Redundant sensing and self-consistency can serve as long-term proxies, validated by higher-accuracy but shorter-term techniques such as surveys, self-reports, video annotation, and controlled experiments.

5.6 Aesthetics Matter in Homes

Myth: *Users won't mind a few sensors around the house.*
Fact: *Aesthetics constrain deployments, especially at large scale and over long time durations.*

Many people will accept sensors into their homes to benefit science, but few want them as decor. Aesthetic appeal is not typically a concern for wireless sensor network design but is important in homes, particularly when hundreds of sensors are deployed over long time durations. For example, our early generation systems were too unsightly to have long-term feasibility. The visual landscape of every room was dominated by wires, exposed circuit boards, and dozens of sensors hanging from the walls, doors, windows, and appliances. The first-generation light switch sensors were particularly obtrusive because they partially blocked the light switch itself (Figure 1(a)). Due to push back from users, we designed our later generation systems to "disappear into the woodwork", quite literally. For example, the height sensors, motion sensors and latch sensors were all fit into a single enclosure that snaps into place behind the door jamb (Figure 1(c)). The COTS light switch sensor hides the electronics behind the switch plate (Figure 1(a)). When designing sensors, choose consistent colors for all components, including enclosures, wires, adapters, tape, and mounting putty. Furthermore, design around enclosures and materials that come in multiple colors, including wood grains if possible. Decrease the visibility of surface-mounted sensors by placing them to maximize balance, alignment and symmetry with the surrounding windows, trim, and other objects.

Leave No Trace: In a residential environment, sensors must not just meet aesthetic criteria when mounted, but also when taken down. Unlike a lab or even an office building where nails and staples may be acceptable, the materials, surfaces, and finishes in homes are often highly refined. We used generous quantities of double-sided tape in our early deployments, but soon found that it peels paint and even plaster from the walls when removed. This is expensive and time consuming to repair, particularly in homes with a different paint color in every room. Our later systems used painter's tape (Figure 1(c)), but even that peels paint when left too long. Our next generation deployments used mounting putty, which works for short-term deployments, but overnight temperature changes will cause it to slowly harden and, in one house with hundreds of putty-mounted sensors, eventually lead to a cacophony of sensors crashing to the ground every 4-6 hours. So far, we found that stretch-release mounting strips by the name brand "3M Command" are the only solution that does not require extensive cleanup and repair once the deployment is over. In addition to paint and plaster, sensors can also cause other surface damage: water near sinks will probably not break sensors but does lead to rust stains over time, and sensors placed on unfinished wood can cause uneven fading from the sun.

No LEDs at Night : In residential deployments, turning off LEDs is not just a power consideration; it is also an aesthetic consideration, particularly at very large scale. LEDs are often considered an annoyance on home electrics devices

such as televisions and alarm clocks, but become a first-class problem when scaling to hundreds of sensor nodes. LEDs on our sensors were barely visible when deploying during the day, but some users complained that the LEDs were so numerous and so bright that it was pointless to turn the lights out, and that the house became a "circus" or a "laser light show" at night. We calculate that our sensors introduced over 150 new LEDs into one home. When deploying at such scale, *all* LEDs must be covered or turned off. It is even necessary to disable LEDs that are placed into plastic enclosures (Figure 1(c)) or deep inside air ducts (Figure 1(e)) because indirect reflections cause these features to glow at night. The short-duration, event-triggered LEDs on our light and faucet fixture sensors (Figure 1(h)) were less of an aesthetic concern, but some users reported that they affected user behavior by making them more cognizant of electricity and water usage, which can be a concern for the scientific validity of some experiments. Even LEDs on outdoor devices should be disabled to avoid the curiosity of animals or passers by.

No Noise: Devices that make noise are a direct risk to continuous data collection and a "deal breaker" for long-term deployments. For example, the ultrasonic transducers in our first generation height sensors (Figure 1(c)) made a constant clicking sound during operation, causing data loss at night when users were forced to disable them. Noise from air resistance in our second-generation active registers was tolerable for the short term, but caused enough noise that they were removed after a few months. Be sure to check whether each device *could* make noise, even if it typically does not. For example, the occupants of one deployment disconnected all data collection devices when an uninterruptable power supply (UPS) began beeping incessantly at 5:30am due to a power outage. Several days of data were lost, and the system would actually have been brought back on-line much faster if the UPS had never been installed in the first place. Some devices such as our 120VAC-to-5VDC converters (Figure 5(a)) cause a high-pitched ringing sound that will cause users to unplug the sensors. These sounds were the hardest to prevent because only a small fraction of devices make the noise due to the high manufacturing variation for cheap electronics, and only some people can hear it due to the very high frequency.

5.7 Simplify the Architecture

Myth: *Industry has already produced a wide range of suitable residential sensing systems.*
Fact: *Many COTS devices were not designed for large scale deployments, and integration of many COTS platforms increases the possible modes of system failure.*

Most large-scale outdoor sensing deployments to date have used custom designs that integrate all sensors into a single system using one wireless bridge, one well-tested software stack, one power source, and a small number of subsystems that need fail safes and redundancy. In the residential domain, sensing and home automation are old industries with a wealth of commercial off the shelf (COTS) products, and in our deployments we used COTS devices whenever possible (Figure 1). These products reduce design time, but complicate the system architecture by introducing multiple protocols and communication paths, multiple bridges and gateways, and multiple software stacks. Deploying a dozen different COTS products in a single home is akin to maintaining a dozen sensor networks simultaneously, each with independent failures, and we found that maintenance effort increases with the number of different sub-systems more than it does with the total number of sensors.

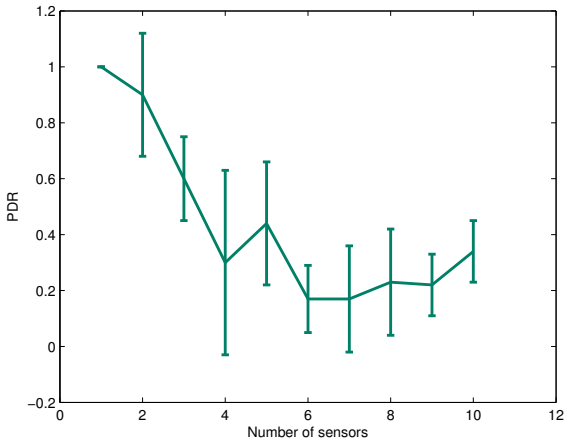

Figure 7. As an increasing number of X10 motion sensors are added to the network, packet deliver ratio (PDR) drops from 100% to near 20%.

For example, in our system, some sensor data was collected by the gateway through a wireless bridge while other sensor data was collected through the home WiFi network and router. Therefore, data is lost if *either* the gateway or the router fails. Furthermore, every new hardware configuration necessitates new robustness mechanisms. For example, our data buffering tools on the gateway provide robustness to broadband connection failure, but not for COTS devices that send data directly to the vendor's server. Similarly, each sensor system that requires a new wireless bridge also necessitates a new software stack on the gateway machine to read from that bridge, which multiplies effort for software development, testing, and on-line maintenance. Heterogeneity of COTS products also increases deployment time because all sensors cannot be prepared and deployed using efficient batch or pipeline operations, such as those used for large-scale outdoor deployments of homogeneous networks [18]. Each additional COTS product introduces another set of tasks, checklists, and risks.

Additionally, most inexpensive COTS products are designed for the hobbyist or enthusiast deploying small-scale home automation or security systems; they are not designed for large-scale deployment or scientific validation. For example, none of the COTS devices have packet sequence numbers, making it difficult to analyze data loss, and some devices such as the TED power meter actually hide data corruption by repeating the last valid value. The GE light switch sensors (Figure 1(a)) do not have reliable transmis-

sion, and must be polled periodically, which increases jitter on event timestamps. Furthermore, these devices are designed to be installed in small numbers, and contain numerous defects that make them very time consuming to install at large scale, including the need for a neutral wire in the gang box; the need to reverse-engineer all 3-way and 4-way switches in the house; wire nuts that make it difficult to fit more than one sensor in a gang box; and metal tabs that must be broken off when installing more than one sensor in the same gang box. These defects added an entire day to the installation time for each house. Wireless devices manufactured by X10 (Figures 1(d) and 1(a)) are designed for very low traffic rates and therefore use a simple wireless protocol with no media access, reliability, or error detection mechanisms: when data needs to be sent it is simply transmitted five times. When deployed at large scale, the packet delivery ratio (PDR) quickly degrades and false data and node IDs begin to appear due to packet corruption. We quantified this performance in a controlled experiment with an increasing number of sensors, using five trials per configuration (Figure 7) and showed that PDR drops to near 20% with as few as 4 nodes in the same radio cell. Thus, COTS devices are a mixed blessing: they shorten design time but increase integration, deployment, and maintenance time, and they are not designed for scientific validation.

6 Conclusions

In our experiences deploying over 1200 sensors in over 20 homes, we observed numerous facts and insights about the challenges, hazards, and pitfalls of residential sensing. These experiences are key for success in future deployments, and in this paper we analyze the full set of all deployments to distill a single set of guidelines, each supported by data, images, and anecdotes from our experience, with the hope of preventing others from repeating our mistakes.

In many ways, the realization that residential sensing is more difficult than just plugging in a sensor and using WiFi is analogous to the earlier realization that deploying a wireless sensor network is more difficult than installing sensors in a lab or testbed. Indeed, many of the challenges of residential sensing are also similar to those first identified for outdoor environments: scale, realities of a hazardous environment, limited access, energy management, and wireless communication. We acknowledge these similarities and embrace them, and have tried to articulate the challenges of residential sensing in those same terms. By doing so, we are also better able to articulate the differences: the nature of limited access, the types of hazards in the environment, and the reasons why large scale and long-duration deployments become challenging.

Acknowledgments

We express great gratitude to the many volunteers in our sensing studies. Thanks to Prabal Dutta for helping interpret the indoor solar data. This work is based upon work supported, in part, by the National Science Foundation under Grants EFRI-1038271, CAREER-0845761, CSR-0720640, and ECCS-0901686.

7 References

[1] Martin Building Group. Improve your health. `http://martinbuildinggroup.com/pdf/health.pdf`.

[2] G. Barrenetxea, F. Ingelrest, G. Schaefer, and M. Vetterli. The hitchhiker's guide to successful wireless sensor network deployments. In *Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 43–56, 2008.

[3] M.C. Mozer. The neural network house: An environment that adapts to its inhabitants. In *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligent Environments*, pages 110–114, 1998.

[4] C.D. Kidd, R. Orr, G.D. Abowd, C.G. Atkeson, I.A. Essa, B. MacIntyre, E. Mynatt, T.E. Starner, W. Newstetter, et al. The aware home: A living laboratory for ubiquitous computing research. *Lecture notes in computer science*, pages 191–198, 1999.

[5] D.J. Cook, M. Youngblood, E.O. Heierman III, K. Gopalratnam, S. Rao, A. Litvin, and F. Khawaja. MavHome: An agent-based smart home. In *Pervasive Computing and Communications (PerCom)*, pages 521–524. IEEE, 2003.

[6] S.S. Intille, K. Larson, J. Beaudin, E.M. Tapia, P. Kaushik, J. Nawyn, and TJ McLeish. The PlaceLab: A live-in laboratory for pervasive computing research (video). In *Proceedings of PERVASIVE 2005 Video Program*, May 2005.

[7] E.M. Tapia, S.S. Intille, and K. Larson. Activity Recognition in the Home Using Simple and Ubiquitous Sensors. *Pervasive Computing*, pages 158–175, 2004.

[8] DJ Cook and M. Schmitter-Edgecombe. Assessing the Quality of Activities in a Smart Environment. *Methods of information in medicine*, 48(5):480, 2009.

[9] Tim van Kasteren, Athanasios Noulas, Gwenn Englebienne, and Ben Kröse. Accurate Activity Recognition in a Home Setting. In *The International Conference on Ubiquitous Computing (UbiComp)*, 2008.

[10] W. Edwards and R. Grinter. At home with ubiquitous computing: Seven challenges. In *Ubicomp 2001: Ubiquitous Computing*, pages 256–272. Springer, 2001.

[11] S. Davidoff, M. Lee, C. Yiu, J. Zimmerman, and A. Dey. Principles of smart home control. *UbiComp 2006: Ubiquitous Computing*, pages 19–34, 2006.

[12] V. Srinivasan, J. Stankovic, and K. Whitehouse. Protecting your Daily In-Home Activity Information from a Wireless Snooping Attack. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 202–211. ACM New York, NY, USA, 2008.

[13] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. *ACM Sensys*, 2010.

[14] V. Srinivasan, J. Stankovic, and K. Whitehouse. Using Height Sensors for Biometric Identification in Multi-resident Homes. In *Pervasive*, 2010.

[15] J. Lu, D. Birru, and K. Whitehouse. Using simple light sensors to achieve smart daylight harvesting. In *The ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, 2010.

[16] Konrad Lorincz and Matt Welsh. MoteTrack: A Robust, Decentralized Approach to RF-Based Location Tracking. In *LoCA*, volume 3479 of *Lecture Notes in Computer Science*, pages 63–82. Springer, 2005.

[17] N. Ramanathan, K. Chang, R. Kapur, L. Girod, E. Kohler, and D. Estrin. Sympathy for the sensor network debugger. In *The International Conference on Embedded Networked Sensor Systems*, 2005.

[18] P. Dutta, J. Hui, J. Jeong, S. Kim, C. Sharp, J. Taneja, G. Tolle, K. Whitehouse, and D. Culler. Trio: enabling sustainable and scalable outdoor wireless sensor network deployments. In *The International Conference on Information Processing in Sensor Networks*, 2006.

[19] J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. *Pervasive Computing*, 2006.