

# A System for Tracking People in Homes for Smart Home Applications

---

A Dissertation

Presented to

the Faculty of the School of Engineering and Applied Science

University of Virginia

---

In Partial Fulfillment

of the requirements for the Degree

Doctor of Philosophy (Computer Science)

by

Timothy W. Hnat

August 2012



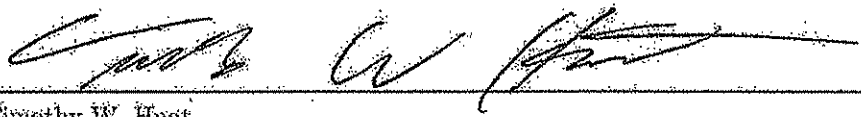
# Abstract

People spend 62 percent of their time within the confines of their home. However, localization technologies such as GPS fail to accurately identify their indoor location. A key requirement of creating a *smart home* is both identifying each person and their current room location. Many techniques for *identification* or *localization* of people in indoor environments are being developed but most require either a privacy-invading camera system or participation by the tracked individuals to either perform actions or wear a device. These systems are useful in an office/industrial setting where the expectation of privacy is different, and an identification badge can be carried. In home environments, active participation and intrusive technologies are not ideal and should be avoided.

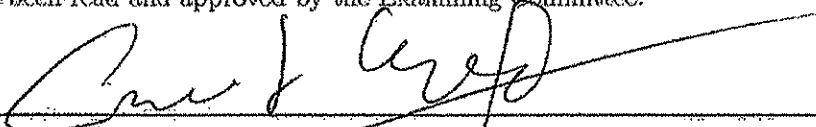
The goals of this dissertation are 1) to identify simultaneously and track multiple people within home environments in a non-intrusive manner and 2) to enable a new set of applications based on location and identity awareness. A custom hardware platform is developed and deployed at the top of doorways to collect both the height and direction of a person as he/she crosses through it. Next, signal processing algorithms are developed to convert the raw data into person-events. Finally, a tracking algorithm, based on simultaneously checking multiple paths, assigns identities and locations to the event stream. The results show the tracking system obtains a 90 percent room-level accuracy on average when tested on 3,000 manually recorded doorway crossings.


## Approval Sheet

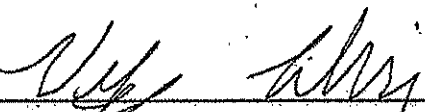
This dissertation is submitted in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy (Computer Science)


  
Timothy W. Hnat

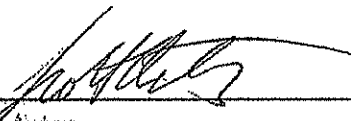
This dissertation has been read and approved by the Examining Committee:

  
Kamm Whitehouse, Adviser

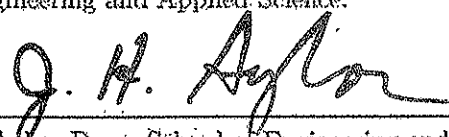
  
Jack Stankovic, Committee Chair

  
Westley Wehner

  
Kevin Skadron

  
Scott Acton

Accepted for the School of Engineering and Applied Science:

  
James H. Ayler, Dean, School of Engineering and Applied Science

August 2012



*To my wife and son*

# Acknowledgements

Thanks to all who provided intellectual, technical, and practical help with various components of this work. These include but are not limited to Tamim Sookoor, Raymond Dawson, Jiakang Lu, Vijay Srinivasan, Erin Griffiths, and my dissertation committee.

My greatest thanks go to my wife, Nastasya, for her constant support.

# Contents

|   |           |
|---|-----------|
| <b>Contents</b>   | <b>v</b>  |
| List of Tables . . . . .                                  | vii       |
| List of Figures . . . . .                                 | viii      |
| <b>1 Introduction</b>                                     | <b>1</b>  |
| 1.1 Problem Overview . . . . .                            | 2         |
| 1.2 Solution Overview . . . . .                           | 2         |
| <b>2 Related Work</b>                                     | <b>4</b>  |
| 2.1 Identification . . . . .                              | 4         |
| 2.1.1 Radio Frequency Identification . . . . .            | 5         |
| 2.1.2 Strong Biometrics . . . . .                         | 5         |
| 2.1.3 Weak Biometrics . . . . .                           | 6         |
| 2.2 Tracking and Localization . . . . .                   | 8         |
| 2.2.1 Binary Sensors . . . . .                            | 8         |
| 2.2.2 Infrared Tracking Systems . . . . .                 | 9         |
| 2.2.3 Acoustic Tracking Systems . . . . .                 | 9         |
| 2.2.4 Vision Tracking Systems . . . . .                   | 11        |
| 2.2.5 Radio Tracking and Localization Systems . . . . .   | 12        |
| 2.2.6 Cellular Phone Localization . . . . .               | 13        |
| <b>3 Hardware Design</b>                                  | <b>15</b> |
| 3.1 Hardware Design Goals . . . . .                       | 17        |
| 3.1.1 Achieving Doorway Coverage . . . . .                | 18        |
| 3.1.2 Power characteristics . . . . .                     | 20        |
| 3.1.3 Cost . . . . .                                      | 20        |
| 3.2 Hardware Design and Operation . . . . .               | 21        |
| 3.3 Challenges . . . . .                                  | 22        |
| 3.4 Early Prototypes and Lessons Learned . . . . .        | 24        |
| 3.4.1 Version 1.0 . . . . .                               | 24        |
| 3.4.2 Version 2.0 . . . . .                               | 26        |
| 3.4.3 Version 3.0 . . . . .                               | 28        |
| 3.5 Alternative Ranging Technologies Considered . . . . . | 30        |
| 3.5.1 Laser Range Finding . . . . .                       | 30        |
| 3.5.2 Infrared Range Finding . . . . .                    | 33        |
| 3.5.3 Infrared Point Cloud . . . . .                      | 35        |
| 3.5.4 Electrostatic Ultrasound . . . . .                  | 36        |
| 3.5.5 Piezoelectric Ultrasound . . . . .                  | 37        |
| 3.6 Evaluation . . . . .                                  | 38        |
| 3.6.1 Platform evaluation . . . . .                       | 38        |
| 3.6.2 Ultrasound Platform Power Analysis . . . . .        | 40        |
| 3.6.3 Cost analysis . . . . .                             | 41        |
| 3.7 Conclusion . . . . .                                  | 42        |

|          |  |            |
|----------|--|------------|
| <b>4</b> | <b>Signal Processing</b>                             | <b>43</b>  |
| 4.1      | Challenges . . . . .                                 | 44         |
| 4.1.1    | Environmental Interference . . . . .                 | 44         |
| 4.1.2    | Measurement Error . . . . .                          | 46         |
| 4.2      | Approach . . . . .                                   | 48         |
| 4.2.1    | Doorway Crossing Detection . . . . .                 | 49         |
| 4.2.2    | Noise Filtering . . . . .                            | 51         |
| 4.2.3    | Height Estimation . . . . .                          | 52         |
| 4.2.4    | Direction Estimation . . . . .                       | 53         |
| 4.3      | Evaluation . . . . .                                 | 54         |
| 4.3.1    | Experimental Setup . . . . .                         | 54         |
| 4.3.2    | Results . . . . .                                    | 56         |
| 4.3.3    | False Detections and Missed Detections . . . . .     | 57         |
| 4.3.4    | Height Measurement Accuracy . . . . .                | 57         |
| 4.3.5    | Direction Measurement Accuracy . . . . .             | 58         |
| 4.4      | Conclusion . . . . .                                 | 59         |
| <b>5</b> | <b>Tracking</b>                                      | <b>60</b>  |
| 5.1      | Design Goals . . . . .                               | 60         |
| 5.2      | Approach . . . . .                                   | 61         |
| 5.2.1    | Models of Height, Direction, and Detection . . . . . | 62         |
| 5.2.2    | Creating Tracks . . . . .                            | 63         |
| 5.2.3    | Weighting Tracks . . . . .                           | 65         |
| 5.2.4    | Merging and Evicting Hypotheses . . . . .            | 66         |
| 5.3      | Simulation-based Evaluation . . . . .                | 68         |
| 5.3.1    | Experimental Setup . . . . .                         | 68         |
| 5.3.2    | False Positives . . . . .                            | 69         |
| 5.3.3    | False Negatives . . . . .                            | 69         |
| 5.3.4    | Sensor Noise . . . . .                               | 70         |
| 5.3.5    | Occupant Height . . . . .                            | 71         |
| 5.3.6    | Direction . . . . .                                  | 72         |
| 5.4      | Experimental Evaluation . . . . .                    | 73         |
| 5.4.1    | Topology Changes . . . . .                           | 75         |
| 5.4.2    | MHT accuracy . . . . .                               | 76         |
| 5.5      | Conclusion . . . . .                                 | 79         |
| <b>6</b> | <b>Deployment</b>                                    | <b>82</b>  |
| 6.1      | Related Work . . . . .                               | 83         |
| 6.2      | Overview of the Residential Deployments . . . . .    | 87         |
| 6.3      | Reliability and Failure Analysis . . . . .           | 90         |
| 6.4      | A Hitchhiker's Guide . . . . .                       | 94         |
| 6.4.1    | Homes are Not a Power Panacea . . . . .              | 95         |
| 6.4.2    | Homes Have Poor Connectivity . . . . .               | 98         |
| 6.4.3    | Homes are Hazardous Environments . . . . .           | 99         |
| 6.4.4    | Homes are Remote Environments . . . . .              | 102        |
| 6.4.5    | Expect Limited User Participation . . . . .          | 104        |
| 6.4.6    | Aesthetics Matter in Homes . . . . .                 | 106        |
| 6.4.7    | Simplify the Architecture . . . . .                  | 108        |
| 6.5      | Conclusions . . . . .                                | 109        |
| <b>7</b> | <b>Conclusion</b>                                    | <b>111</b> |
| 7.1      | Future Work . . . . .                                | 112        |
|          | <b>Bibliography</b>                                  | <b>114</b> |

# List of Tables

|     |  |    |
|-----|--|----|
| 3.1 | Hardware Evaluation . . . . .            | 38 |
| 3.2 | Power Consumption by Component . . . . . | 40 |
| 3.3 | Bill of Materials . . . . .              | 41 |
| 4.1 | Precision and Recall . . . . .           | 57 |
| 4.2 | Direction Accuracy . . . . .             | 58 |
| 6.1 | Deployment Overview . . . . .            | 86 |
| 6.2 | Sensor Downtime . . . . .                | 91 |
| 6.3 | Total Downtime . . . . .                 | 92 |
| 6.4 | Average and Peak Power . . . . .         | 97 |

# List of Figures

|      |   |    |
|------|---|----|
| 3.1  | Design Illustration . . . . .                               | 16 |
| 3.2  | Sensor Design . . . . .                                     | 17 |
| 3.3  | Illustrations of Design Parameters and Challenges . . . . . | 18 |
| 3.4  | Tilting Sensor . . . . .                                    | 19 |
| 3.5  | Discrete Mounting . . . . .                                 | 20 |
| 3.6  | Door Coverage . . . . .                                     | 23 |
| 3.7  | Go!Motion . . . . .   | 25 |
| 3.8  | Maxbotix . . . . .  | 26 |
| 3.9  | Version 3.0 Sensor . . . . .                                | 28 |
| 3.10 | Version 3.0 Sensor . . . . .                                | 28 |
| 3.11 | Version 3.0 Schematic . . . . .                             | 29 |
| 3.12 | Laser Distance Sensor . . . . .                             | 31 |
| 3.13 | Light Detection and Ranging Sensor . . . . .                | 33 |
| 3.14 | Angle of Reflectance Sensor . . . . .                       | 34 |
| 3.15 | Infrared Point Cloud Sensor . . . . .                       | 35 |
| 3.16 | Go!Motion Electrostatic Sensor . . . . .                    | 36 |
| 3.17 | PING Piezoelectric Sensor . . . . .                         | 37 |
| 4.1  | Environmental Hazards . . . . .                             | 45 |
| 4.2  | Multi-modal Noise . . . . .                                 | 46 |
| 4.3  | Noisy Readings . . . . .                                    | 47 |
| 4.4  | Hand-held Objects . . . . .                                 | 48 |
| 4.5  | Noise Filtering . . . . .                                   | 49 |
| 4.6  | Event Types . . . . .                                       | 50 |
| 4.7  | Example Noise . . . . .                                     | 51 |
| 4.8  | Multi-path Errors . . . . .                                 | 52 |
| 4.9  | Sensor and Doorway Locations . . . . .                      | 55 |
| 4.10 | Bi-partite Matching . . . . .                               | 56 |
| 4.11 | Height Accuracy . . . . .                                   | 58 |
| 5.1  | Ambiguous Events . . . . .                                  | 61 |
| 5.2  | Location Resolution . . . . .                               | 67 |
| 5.3  | False Positive Analysis . . . . .                           | 70 |
| 5.4  | False Negative Analysis . . . . .                           | 71 |
| 5.5  | Sensor Noise Model . . . . .                                | 72 |
| 5.6  | Sensor Noise Analysis . . . . .                             | 73 |
| 5.7  | Sensor Timeout Analysis . . . . .                           | 74 |
| 5.8  | Occupant Height Analysis . . . . .                          | 75 |
| 5.9  | Direction Analysis . . . . .                                | 76 |
| 5.10 | MHT Description . . . . .                                   | 77 |
| 5.11 | Various House Topologies . . . . .                          | 78 |
| 5.12 | Topology Depth Analysis . . . . .                           | 80 |

|      |  |     |
|------|--|-----|
| 5.13 | Tracking Results . . . . .                           | 81  |
| 6.1  | Sensors Deployed . . . . .                           | 88  |
| 6.2  | System Architecture . . . . .                        | 90  |
| 6.3  | Total sensor down time for each deployment . . . . . | 93  |
| 6.4  | Solar Energy . . . . .                               | 95  |
| 6.5  | Hazards . . . . .                                    | 100 |
| 6.6  | System Verification . . . . .                        | 103 |
| 6.7  | X10 Packet Delivery Rate . . . . .                   | 108 |





# Chapter 1

## Introduction

Since man realized that he could travel beyond the borders of his home location, he has desired to know where he is located and how he is going to get to his next destination. The use of a sextant and the stars, a compass and topology map, and the modern Global Positioning System [1] (GPS) have each in their turn revolutionized the ability of man to know where he is. This steady improvement of localization abilities, however, is limited to the outdoors. No technology localizes or tracks people while inside buildings without relying on active participation by the occupants, either by carrying or operating a technological device. This interferes with the process of just living life, and if a better system could be devised, the new and exciting opportunities for future research and applications are endless.

Companies such as Google or Facebook rely on personal information either from what a user searches for or who and what his or her friends and interests are. Personalized searches provide useful information to the person looking up information, thus improving his or her experience with these services. In each case, understanding the habits and interests of the users allow each company to provide a better service. However, these can only exploit digital activities on the Internet. According to a recent survey [2], people spend 62 percent of their time within their homes. If a system existed that could identify daily patterns and activities of people at home, many new and revolutionary applications could be built. For example, the knowledge of which room a person is located in supports individualized activity recognition applications or an automation system that could tune the energy consumption of the house. These are two of a wide range of applications that will be possible.

## 1.1 Problem Overview

Reliably tracking and identifying residents within their homes without active participation leads to the possibility of many exciting applications. For example, elderly people may be able to maintain their independence in their own homes without the necessity of carrying a LifeAlert<sup>TM</sup> device or cellular phone with them at all times. Home health monitoring could improve the treatment of conditions, such as dementia or depression which are difficult to diagnose and monitor with only periodic visits to a doctor. If activities, locations, and energy use can be attributed to individuals within a home, then targeting those specific rooms and activities rather than the whole house could reduce energy consumption as a whole.

Existing technologies require wearable devices or video cameras, which means participants must carry devices or interact with in-place technology disrupting their lives and limiting applicability. Technology can be misplaced or forgotten altogether, thus rendering the system unable to perform its intended task. Placing cameras within participants' homes when deploying and running experiments present the possibility of privacy invasion. Most people do not want nor allow any video recording equipment to be used in their home. These two factors combine to make deploying a real-world system problematic.

## 1.2 Solution Overview

The solution presented in this dissertation removes these two obstacles by proposing to track and identify residents by passively measuring their height as they pass key points in the home such as doorways. Despite substantial measurement noise, use of low-power wireless ultrasound sensors to track multiple people with room-level accuracy performs as well as or better than existing technology without the large privacy invasion or inconvenience.

By using height as a weak biometric, this system avoids the active participation of occupants, allowing them to live their lives normally. It also avoids the privacy invasion of using cameras. Height was chosen because, while not a uniquely identifying feature worldwide, it is typically enough to distinguish people in homes. It does have limitations, but so do other weak biometrics that have been considered in the past; weight fluctuates or is defeated by carrying objects, color recognition requires the use of a camera and clothing changes daily, and gait analysis requires either sensors placed on the participant's body or cameras. Height is fairly reliable and can be measured with a small number of sensors.

Chapter 3 presents a custom-designed hardware platform for sensing height and direction as people walk through doorways. The hardware addresses three aspects of sensing: (1) physical characteristics derived from different properties of the environment and the human subjects; (2) measurement characteristics directly related to the quality of data produced by the sensors; and (3) energy characteristics that reflect the amount of power required to run the devices. Each aspect is composed of several design challenges, all which must be globally optimized to create the final hardware platform responsible for recording and transmitting all data readings to a central location for analysis.

Chapter 4 presents a signal processing algorithm that converts data produced by the hardware platform into a sequence of events. It addresses two goals. First, it detects and records a person walking under the doorway based on raw data while minimizing both false positives, when noisy signals are detected as events, and false negatives, where real events are missed. Second, it determines the height and direction for each event. Environmental interference, environmental change, target interference, and sensor noise are all unique challenges associated with these goals. The signal processing algorithm detects 96 percent of events within a house by addressing these goals and challenges.

Chapter 5 presents a multiple hypothesis tracking algorithm that simultaneously identifies and tracks residents based on the events from signal processing. This indoor tracking technique is an enabling technology for various context-aware applications. The tracking system meets the following two goals: (1) Individual people must be differentiated. (2) People must be localized with room-level accuracy. A simulator is utilized to test and optimize various parameters of the tracking algorithm.

The three hardware designs discussed in this dissertation were built and deployed in four homes over the course of a year and a half. The first deployment lasted only two weeks but served as proof of concept. Through lessons learned from both the first and second hardware designs, the final solution was conceived, deployed, and currently still provides data. These real world tests allow a holistic evaluation because long-term ground truth information is challenging to obtain, a series of short experiments are utilized to test the platform. The system averages 90 percent tracking accuracy across all experiments, even though the participants carry no technology. In comparison, the baseline active tracking system that uses radio frequency beacons, Motetrack, only achieved a 37 percent tracking accuracy.

Chapters 3, 4, and 5 are published as a single paper in *The ACM Conference on Embedded Networked Sensor Systems (Sensys)* 2012. Chapter 6 is published in *Sensys* 2011.

## Chapter 2

# Related Work

Solutions to identifying and localizing individuals and objects have been developed for decades [3,4]. Identification systems associate an identifier with an individual or object. There are several types of identification systems including radio frequency identification, strong biometrics, and weak biometrics which all have varying distinguishing features. Localizing systems are designed to determine an object's or person's position, while tracking measures a changing position over time and disambiguates object identities when paths of multiple objects cross. A variety of identification, localization, and tracking solutions are described below, including systems based on infrared transmitters, pyroelectric sensors, acoustics, vision-based systems, radio-based devices, and cellular phone systems.

### 2.1 Identification

Identification systems are designed to associate identities to individual persons. For example, radio frequency identification tags (RFID) contain a unique number that can be read by antennas that are stationed throughout the monitored area. These tags are carried by a person, and are typically used to control access to various secured environments.

Carrying a tag is preferred in “secure” buildings but in home environments, people tend not to carry any form of technology all the time. In contrast, biometric identification uses features of the person's body and is split into two classes. First, *strong biometrics* such as a fingerprint, iris scan, or facial features uniquely identify a person and typically require an action such as placing a finger on a reader or looking into the iris scanner. Second, *weak biometrics* such as height, weight, or gait can be utilized to help distinguish between individual in a small group of people but cannot uniquely

identify them in a global group; however, these features can usually be measured without forcing a person to perform a particular action by remotely monitoring each person with cameras, force plates, or other non-contact sensors. A wide range of identification system is described below.

### 2.1.1 Radio Frequency Identification

Radio Frequency Identification (RFID) systems offer an excellent way to track and identify objects. Their low-cost tags are easily placed on every object while a powerful reader can read them from a distance. There are two different classes of RFID tags, active and passive.

An active tag contains a power source and is capable of transmitting data over a greater distance. LAND MARC [5] is based on active tags and utilizes triangulation of signal strengths between tags and readers to localize objects. It also utilizes a set of “reference” tags that are placed within the environment as a way to measure RF characteristics.

Active tags allow for a greater range than traditional passive tags, although they are hampered by their size and power supply. In contrast, passive tags have no concerns about power, but their effective range is significantly reduced. Chawla et al. [6] utilize multiple RFID readers and varying power levels to localize these tags within a space.

While RFID tags provide an excellent and low-cost way to identify and localize an object, they are intrusive due to the need to carry a tag around with oneself at all times. In the future, smart clothing may already contain these tags and the problem of carrying tags will effectively disappear, but a new problem of associating tags with people will be introduced. The benefit of the system presented in this dissertation is that there is not a need for a person to carry around any tag, reducing the impact on the person being tracked.

### 2.1.2 Strong Biometrics

Strong biometric systems identify people based on physical characteristics that are universal, unique for each person, measurable, and do not vary through time [7]. Additionally, the process of collecting biometrics must be acceptable to the people being identified. Many strong biometric systems have been developed, and three representative types of biometric identification solutions are discussed here: fingerprint, iris, and facial.

A fingerprint identification system [8] is the most common mechanism for providing a strong biometric for unique identification purposes. Each person is assumed to have a uniquely identifying

fingerprint. Fingerprints are utilized in security scenarios as additional authentication because they are difficult to forge.

Iris patterns provide a unique texture similar to fingerprints that is stable as a person ages and can be used for identification purposes [9]. However, unlike a fingerprint system where the person must touch a reader, iris scanners do not require a conscious action. They do require a camera capable of recording the details in the eye [10] and for the person to look into a camera for the picture to be made. Automated systems that require no user participation are not currently commercially available.

Facial features move further into the realm of remote sensing by capturing an image of the human face [11]. Features are computed and matched against an existing database. Research shows that once a high quality image of the face is captured it can be used to uniquely identify a person; however, this requires a camera system capable of tracking and recording faces of people. This is not practical for smart home applications because multiple cameras would need to be placed at eye-level in each room to reliably capture this data. Additionally, people can move around the home at night in low light, which prevents a camera from capturing the image. Additionally, cameras have the potential to cause privacy concerns and may be rejected by the home occupants.

### **2.1.3 Weak Biometrics**

The goal of using weak biometrics is to provide a less invasive mechanism to identify a person by exploiting physical features that are easy to measure, even if they are not necessarily uniquely identifying. Traditionally, biometric data is difficult to gather because it requires the person to alter his/her actions to authenticate with the biometric reader, or expensive readers that can record from a distance. This can be as simple as a fingerprint scanner or as complex as an iris scanner. The key challenge when utilizing weak biometrics is accommodating their non-unique identifying nature. However, because they are much easier to acquire from a distance, the tracking system relies on measuring physical properties of people. There are a large number of weak biometrics that have been utilized for identification and several representative examples are discussed in this section.

#### **Weight**

The first weak biometric is the weight of a person. Weight is an identifying feature for people if it can be measured accurately and does not change quickly; however, it can easily be defeated by carrying an object. Jenkins et. al. [12] show that by using a load cell, the weight of a person walking over

it can be accurately extracted; however, this information is not sufficient for tracking. The Sensory Floor [13] was built by installing an array of load cells in the floor of a room. Other projects [14–18] utilize similar ideas. The benefit of using weight is that each person does not have to do anything differently or carry any object. The major drawback to using weight as a tracking system is the high-cost of installing load cells to measure the weight of a person as he/she steps on it. There are two techniques for deploying this system. First, thin load cells can be placed on the ground around each doorway so that a person must walk on them to switch rooms. The second approach involves constructing an overlay for the entire floor in the room. The first approach is unsightly and are a potential tripping hazard. The second approach is costly and not practical for entire buildings.

### **Gait**

Gait is the pattern of movement of the limbs while walking or running. Several projects [19–21] have utilized gait analysis to identify different people. Liu et. al. [19] utilize a video camera to capture the gait of a person from a distance and combine it with facial recognition to assign an identity to a person. This technique is limited to what a camera can record with sufficient detail and can fail to function correctly when placed in the confines of a house because of the lack of clear sight lines in each room. Gafurov et. al. [20] explore a different way to capture the gait of a person by attaching low-power sensors to measure the limbs’ motion. They are able to use the extracted gait features to provide periodic verification of a person’s identity for security authentication. The downside to this approach is that various sensors must be placed on a person and that it does not apply directly to the tracking problem. One can construct a tracking system that should perform well by combining multiple techniques such as using the on-board sensors of a cellular phone or RFID-based sensors. However, requiring a person to wear or carry sensors poses problems: these sensors can be removed or forgotten, creating incomplete and/or inaccurate data.

### **Height**

Srinivasan et. al. [22] have proposed using height as a tracking system and show that identity can be obtained through repeated measurements. Several systems were built and deployed in various houses to gather real-world data sets based on this weak biometric height information. However, this project does not implement a tracking system based on this observation and only validates identities.

The challenge with utilizing any of the weak biometric attributes is that they are not unique identifiers; however, as identified by Srinivasan, complete uniqueness is not a requirement for a

tracking application designed to monitor a handful of people inside a house. Weak biometrics allow a system to record data about people as they go about their daily lives and provide enough separation between individuals; however, it has not been demonstrated as a tracking system.

## 2.2 Tracking and Localization

Many would consider localization of a person in an outdoor environment to be solved by the Global Positioning System [1] (GPS), which revolutionized determining location anywhere in the world. In indoor environments, tracking and localization has been extensively studied with many systems such as Active badges [23] that utilize infrared transmitters, pyroelectric [24] systems that measure infrared signatures, Pioneer Robots [25] and Millibots [26] which use acoustic transmitters, the CITRIC [27] system based on a camera system, facial recognition [28] systems, and radio-based technology with Motetrack [29]. Built on a variety of technologies, some have the ability to assign identities while others simply understand where people or objects have moved. There are different technology classes including binary sensors, infrared tracking, acoustic tracking, vision-based tracking, radio tracking and localization and cellular phone systems.

### 2.2.1 Binary Sensors

One class of tracking systems relies entirely on binary sensors, which only produce an on or off state. The most common binary sensor is a motion sensor that two separate projects utilize to perform indoor localization. Both projects use a similar approach for localization through the use of models and clusters of binary motion sensors to isolate a person.

The first system [30] uses a particle filter to combine the binary motion sensor inputs, modeled with conditional probability functions, into a two-dimensional room location. It also relies on a model of the person's motion. The goal is to record the position of a single person within the confines of a single room. The second project [31] also targets tracking people with binary motion sensors but utilizes a centroid-based localization scheme. Centroid localization is equivalent to taking the center of mass for all sensors that see the object. Their research is focused on how to cluster and group the sensors rather than tracking a person.

Both projects target specific problems and needs but do not address what happens when multiple people are present in their environments. Both of these systems have no way to assign an identity to a person. The final solution and future applications rely on knowing the identity of each person within the house.



### 2.2.2 Infrared Tracking Systems

A separate class of tracking systems utilizes infrared (IR) transmitters as a localization technique. Active badges [23] are small devices that are worn by people or placed on objects, which transmit an encoded IR signal. This signal is picked up by receivers within the environment and is limited to about 10 meters. One concern with IR signals is that they are typically line-of-sight signals. Anytime the transmitter or receiver is obstructed its ability to communicate is lost. Fortunately this is typically not a problem within a house since signals can reflect off most items in a room thereby creating a line-of-site path. As long as the transmitter and receiver are not obscured in a way that removes all possible reflection paths between the pair, they will function properly.

Utilizing badges for tracking applications is a problem even though they only weigh 40g. Badge systems function well in a work environment where they are typically worn around the neck or clipped to clothing and usually contain an identification picture. This is minimally invasive and perfectly acceptable for corporate security. However, people will only tolerate and wear this type of device within their home for a limited time because clothing changes and different clothes may or may not have a place to attach the device.

#### Pyroelectric

Pyroelectric sensors measure the environment to detect information about its occupants. Simple versions of the sensors are known as Passive Infrared (PIR) sensors and only measure movement. Pyroelectric sensors additionally measure angular velocities. Hao et. al. [24] built a tracking system consisting of multiple pyroelectric sensors and used the intersections of the angular measurements to triangulate where a subject is within a room. This system relies on maintaining target information to disambiguate multiple people; however, it does not assign a true identity to each person. Multiple room tracking needs to be able to maintain independent identities without needing to keep specific state information.

### 2.2.3 Acoustic Tracking Systems

A variety of systems employ acoustic systems for both device localization and object tracking. There are two separate classes of acoustic localization techniques: transmitter-oriented where the tracked objects produce the sound and receiver-oriented where the objects listen; however, both rely on the Time-of-Flight (TOF) of sound through the air, which is approximately one foot per millisecond. First, acoustics are used to locate the nodes within the sensor networks. Then, acoustics are used

to track and localize people/robots in their environments. Pioneer Robots [25] are one of many projects that employ active acoustical localization to identify the location of a robot. Millibots [26] utilize a set of small autonomous robots that contain ultrasonic beacons. An ultrasonic and radio pulse are sent concurrently to other robots. The RF pulse travels at the speed of light and will arrive at all receivers at the same instant while the sound wave will lag. The robots can determine the distance from the source by timing the lag. Once all timings are known for the robot set, this information is fed to a central computer, which uses trilateration to determine the position of all robots. Cricket [32] employs a nearly identical technique for localizing a mobile person except that the beacon nodes are placed carefully within each room. The listening device a person carries picks up the beacons to determine its location as he/she walks around the environment.

The Medusa Kindergarten [33], AHloS [34], ENSBox [35], and a project from UIUC [36] are different platforms that utilize ultrasonic collaborative multilateration to localize a network of nodes with only a few known beacon locations. AHloS provided fine-grained localization for sensor networks and forms the basis for *multilateration*, a process for locating an object by computing the difference in arrival times of two signals from three or more receivers. The UIUC project utilizes consistency checks and statistical signal processing to help avoid problems with noisy signals. ENSBox utilizes an array of microphones and a high-power CPU to precisely sample acoustic pulses. The multiple microphone array allows the bearing to be determined and identifies the direction of the sound source, further aiding localization. Distributed Object Locating System for Physical-space Internetworking (Dolphin) [37, 38] follows a similar philosophy and minimizes the number of beacon nodes necessary. Nodes localize themselves and then use this information to localize others by through multi-hop communication.

Active Bats [39] function in a similar manner as Millibots or Cricket; however, they differ in how the ultrasound is utilized. Instead of computing the location of a device on the mobile node by listening to existing beacons, this system reverses the process. The Bats actively transmit ultrasonic pulses, which are picked up by receivers in the environment. Each tag transmits an encoded signal to uniquely identify itself to the system, thus supporting multiple users.

LaSLAT [40] or *Laplace Simultaneous Localization and Tracking* is a technique that relies on a target to emit periodically ultrasonic events as it is moved through an environment. It utilizes these same pulses to both localize and calibrate the in-place network, thus simultaneously localizing and tracking. This process speeds up the deployment time, and tracking accuracy improves over time as more measurements occur. Acoustic localization systems rely on a device that is carried by each object being tracked which can be forgotten or misplaced.

### 2.2.4 Vision Tracking Systems

Many researchers have utilized various image processing techniques in conjunction with video cameras to construct tracking systems. Three separate projects represent various vision-based tracking systems.

CITRIC [27] is a camera system that utilizes the Intel XScale processor and a Tmote Sky mote for communication. People are tracked through a scene by utilizing a simple image segmentation algorithm. Multiple targets can be tracked simultaneously as long as the subjects maintain separation. CITRIC is appealing because all difficult computation is performed locally on each node, and a low-dimension feature set is sent to a central location to perform the tracking algorithm. There are many other projects that utilize a camera and mote platform [41–45] to perform localization.

There is a limitation on how accurate localization can be when only utilizing a basic camera. It is challenging to tell the distance between objects and the camera without making some assumptions about their dimensions. Stereo cameras solve this problem by placing two cameras a known distance apart and utilizing the slight differences between the images to compute distance through triangulation. Once distance is known, the spatial location is much easier to compute from the two coordinates plus distance. Two projects use this exact technique to localize people within environments [46, 47] though both suffer from the same problem. The maintenance of persistent identity through a long-term application scenario is difficult. Vision-based techniques rely on color information to segment and identify the individuals. This is done by computing a histogram of the colors for a segmentation and using that as the identity and works well for tracking objects within the context of a few cameras. However, it breaks down once the clothing a person is wearing changes or people wear similar colors. The hardware system utilized in this dissertation is color independent because no sensor was deployed to measure this information. Instead, other information is utilized that is more difficult to change quickly such as height or weight.

A number of challenges arise when using a video-based localization system. First, cameras are limited to what they can see. Anytime something is obscured from their view such as around a corner or behind a box, then no information is known about what is happening. This is quite different when compared to the various radio-based approaches, which easily go through obstacles and walls. Second, while video surveillance systems are acceptable in most areas around the world in public places and work environments, they are not well accepted in private settings such as a home. Most people object to having any camera system unless they can fully control it, like a baby monitor. This makes deploying a camera-based system difficult from a social acceptance point of view. Finally,

most vision systems suffer from an inability to assign a persistent identity to individuals across multiple interactions. Without a consistent identity, the usefulness of localization for smart home applications is significantly reduced.

The final hardware platform is built to avoid social challenges associated with video systems by measuring information that does not violate personal privacy. Second, it utilizes weak biometrics, which allow it to maintain persistent identities for long-duration applications and experiments.

### 2.2.5 Radio Tracking and Localization Systems

One of the earliest indoor radio tracking systems is RADAR [48]. Developed at Microsoft, this system utilizes multiple base station triangulation techniques that involve a mobile radio device that broadcasts beacon messages to multiple receivers placed throughout the area. This system was built on existing wireless cards and utilized the signal strength and signal-to-noise ratios. Multiple nearest-neighbor receivers are utilized to compute the weighted center of mass resulting in error of approximately 3.5 *m*.

Motetrack [29] introduces a signature-based approach to the radio tracking problem. The system is first calibrated with a set of broadcasting beacon nodes and a receiver that is moved throughout the environment collecting data. Once the signatures are created, they are used to reproduce locations based on the signal strength and which base stations are heard. This signature-based approach produces a 2-3 *m* accuracy throughout the building. The downside with utilizing the lower power transmitters that Motetrack is based on is the necessity to include a large number to keep these accuracy numbers.

RADAR and Motetrack are limited to two-dimensional locations. Liang et. al. [49] take a similar approach by using beacon and receiver nodes to locate a person; however, they restrict all beacon nodes to be on the ground floor and deployed in a non-symmetrical manner. This constraint allows localization to occur in three dimensions without having to worry about how many nodes can be heard by the receiver. The localization is built from the ground up, utilizes the standard *Received Signal Strength Indicator* (RSSI), and is limited to a certain height.

MERIT [50], *MEsh of RF sensors for Indoor Tracking*, addresses one of the primary concerns with all existing radio-based tracking systems up to this point. The authors show that a site profile or site survey is not necessary to obtain accurate localization information. MERIT also identifies that obtaining the room that a person is in is sufficient instead of the true (X, Y, Z) coordinates. The problem is that RF signals easily travel through most walls, which make them excellent for

applications without line-of-sight but present a problem when looking to place a person in a room. The main idea here is to create an RF shield, which keeps RF energy directed into a specific room. This avoids multi-path reflection problems and allows the standard RSSI metric to be used. The accuracy of this system approaches 99 percent for room-level localization.

*Ultra-Wide Band* (UWB) [51] is a radio technology that transmits on a large frequency band. Its key difference from traditional narrow-band radios is the lack of reflected signals. Because there are so many frequencies, some of these will travel directly to the receiver without reflecting off or being absorbed by any surfaces. This allows a more precise indoor tracking system to be constructed. The average error for a tracking system built with UWB radios is approximately 25 *cm*. This is an order-of-magnitude better than other techniques.

Radio-based tracking systems provide a variety of solutions to the problem of localizing a person in an indoor environment. The major drawback with all these techniques is a requirement that a powered device, either a receiver or transmitter, be carried on the person's body at all times. This class of systems may be acceptable in a corporate/work environment; however, in a home setting devices will be forgotten and misplaced unless attached to the body. Another drawback is that the older narrow-band radio technology can only achieve an accuracy of up to 2 *m*.

### 2.2.6 Cellular Phone Localization

Many researchers have proposed using cellular phones to localize people. This makes a lot of sense on the surface. First, almost everyone carries a cellular phone with them during the day. By taking advantage of an existing technology, they hope to gain the tracking ability with minimal overhead. Cellular providers already offer a service for localizing phones as part of the E-911 service designed for emergency response as an alternative to GPS. However, precision is limited to 300 *m*.

The CILoS [52] project attempts to improve on indoor cellular phone localization by measuring RSSI signatures from nearby cellular towers and using these signatures to identify the phone's location. They report an average localization error of 5 *m*. While much more precise than the 300 *m* required for emergency services, it is not precise enough. The 5 *m* accuracy will identify which part of an industrial plant a person is located, but it is not enough to reliably localize a person to a specific room in a residential building.

Another approach [53] utilizes the GSM cellular network to localize people in indoor environments. This approach uses more towers than a phone would use for communication. While many visible towers are not close enough for communication, their signal strengths can be used to build a dense

fingerprint to localize people. This increased set of measurements results in an accuracy of about 2.5 *m* within a single floor of a structure and with sufficient vertical resolution to determine on which floor a person is located.

SurroundSense [54] is a completely different approach to localizing a cellular phone. This approach uses the suite of on-board sensors to record other information about the environment in which it is located. The ambient audio is recorded then matched to previously visited locations, such as department stores. Light and Wi-Fi readings are also recorded with the audio to form a fingerprint, which yields an 87 percent accurate localization. Audio fingerprinting will not work well within the confines of a home environment. There are two reasons for this: (1) many rooms will have nearly identical quiet signatures and (2) signatures for individual rooms could vary widely.

Approaches that utilize cellular phones seem like a logical choice for most localization problems. These high-power, sensor-rich devices are nearly ubiquitous in the modern world and excellent choices for commercial buildings. People will carry cellular phones everywhere while they are away from home; however, once they enter a house, this assumption breaks down. Many times a person will place objects carried throughout the day on a table or shelf, which limits how they can be utilized indoors. Since a cellular phone is unlikely to be carried around reliably, another technique is necessary.

The research community has produced a large variety of tracking, localization, and identification systems; however, none combine to form a complete solution without requiring an action from the person to be tracked or forcing them to allow video recording within their homes. This dissertation shows how a system can be constructed that requires no significant action, such as carrying a device or altering their behavior. By utilizing carefully chosen hardware, weak biometric measurements are used to identify, localize, and track each person within an instrumented house.

## Chapter 3

# Hardware Design

The vision that motivates this work is a sensing device capable of running on battery power for six months that measures both the height and direction of a person walking through each doorway (Figure 3.1). This platform (Figure 3.2) is composed of a sensor suite providing it with the ability to measure height under the doorway, motion on either side of the doorway, and whether the door is open or closed. Some additional environmental sensors, such as temperature, humidity, light, and magnetic, can be included to facilitate and augment future applications but are outside the scope of the analysis in this chapter.

Hardware design is driven by a series of requirements that fall into three distinct classes: (1) physical characteristics of the environment and the human subjects being measured; (2) measurement characteristics of the sensors doing the measurement; and (3) energy characteristics of the entire hardware platform.

Several ranging technologies were considered to solve this problem. For example, laser-based systems rely on timing the transmission and reflection of a laser beam when it reflects off a target. Light detection and ranging (LIDAR) based systems utilize lasers to produce three-dimensional models of the environment. Point cloud and infrared ranging systems rely on measuring the reflection angle of an infrared dot to determine distance. Each of these offers a different trade-off in terms of the questions that must be answered when designing this hardware platform:

- How much will the system cost?
- How many sensors are needed?
- How much energy does the device consume?




Figure 3.1: The goal of the hardware design is to record the height and direction of a person as he/she walks through the doorway.

- What is the sensing range?
- How fast does the system need to sample?

A global optimization of all design parameters is required to achieve the best possible design, which is difficult due to the discrete and discontinuous nature of the hardware design space. Instead, several different devices were built and tested in a lab setting with three different devices deployed and tested in houses. First, a high-powered, narrow beam electrostatic ultrasound device was deployed at the center of each doorway for a two week trial period where data was collected through long USB cables. Second, a device consisting of a wide-angle, low power piezoelectric ultrasound sensor and wireless transmitter was deployed for six months in a single house. The final solution consists of multiple narrow angle piezoelectric ultrasound sensors on each doorway. This system has been deployed for approximately 18 months in four different houses for testing and development purposes. This chapter is published as part of a *Sensys* 2012 paper.






Figure 3.2: A model of a hardware platform that can be placed into a doorway and appear to belong there.

### 3.1 Hardware Design Goals

The primary goal of the hardware platform is to measure the heights and directions of people as they walk through a doorway. Secondary goals of the platform include cost, energy consumption, and a discrete and non-obtrusive aesthetic. In order to be cost effective for smart home applications such as energy management, the system should cost less than a few hundred dollars per home. A typical home has 10-15 doorways including exterior doorways so the components of the system should total \$20-30 at production scales. The devices must also be able to operate for multiple years on battery power: although power is abundant in homes, it is typically not available above the doorways. Running wires from the top of each doorway would increase installation cost and cause a snagging hazard. Additionally, the sensor will be placed in a highly visible part of the home and must have a discrete design and a form factor that can be easily mounted.

Recording when a person walks under a doorway sounds simple, but several conflicting goals must be addressed when designing a hardware platform. Doorways and the way people walk through them vary widely, and robust features must be extracted. Furthermore, each sensor has a different noise and power profile. These three characteristic sets - physical features, measurement capabilities, and power - combine to lay out the desired properties of the hardware platform.

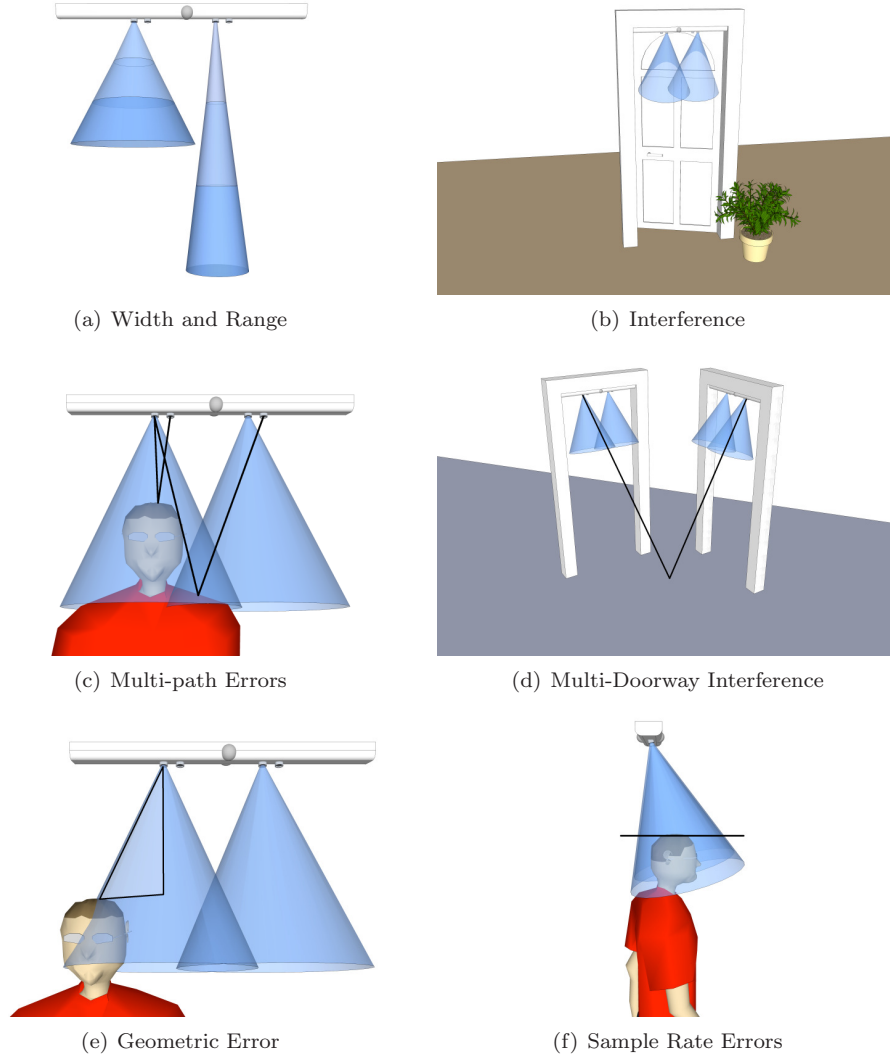


Figure 3.3: Illustrations of different design parameters and challenges

### 3.1.1 Achieving Doorway Coverage

A doorway sensor must operate under a wide range of doorway sizes, person sizes, and walking speed, while still achieving high height and direction accuracy, few false detections, and few missed detections. Doorway sensors must provide 1 *cm* resolution for people with heights ranging from 151 *cm* (5 ft) to 189 *cm* (6 ft, 2 in)<sup>1</sup>, at walking speeds of up to 3 m/sec<sup>2</sup>, in doorways that range from 90-300 *cm* (3-10 ft) wide and 213-275 *cm* (7-9 ft) tall. In order to detect a person, the sensing region must cover the entire width of the doorway and the sensor must sample quickly enough to ensure at least one reading to the top of a person's head when the person is directly under the doorway. The sensing region must also extend outside the door frame in order to detect the direction

<sup>1</sup>The 5th and 95th percentiles of adult human height [55].

<sup>2</sup>More than twice the average adult walking speed [56].




Figure 3.4: The ultrasonic sensor is tilted at approximately 15 degrees from vertical away from the doorway. This creates an asymmetric sensing region that allows the sensor to determine both height and direction.

from which the user enters or exits the doorway. At the same time, the sensing region must not extend too far from the door frame to avoid erroneous measurements such as a person walking by the doorway but not through it.

Finer resolution than 1 *cm* is not expected to help due to a natural variation in posture and gait. It must also be able to sample quickly enough to get an accurate reading in the short time period when a person passes under the door frame. The hair, skin, and clothing on people are *soft targets* that can deflect and absorb a ranging signal, instead of reflecting it back to the sensor. Signals can also be deflected away from a receiver by the curved shape of the head or other parts of the body that are not perpendicular to the door frame. This makes ranging in doorways more error prone than many industrial settings where smooth, hard surfaces on machinery provide clear reflections.



Figure 3.5: The doorway sensor is discretely mounted behind the doorjamb in a molded plastic enclosure.

### 3.1.2 Power characteristics

While power is available in many locations within a house, it is typically not above each doorway. Therefore, power from wall outlets is typically difficult to route to the top of every doorway and potentially inconveniencing the occupants of the house, violating one of the design principles.

The amount of *energy* a sensor uses is of great concern when deploying in locations without easily accessible AC wall power such as a doorway. Some sensor hardware consumes significant power and might not be possible to run off of batteries or use energy harvesting techniques such as a  $6 \times 9$  cm solar cell placed within home environments. 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 [57].

### 3.1.3 Cost

The cost of the system has to be justified by the value of the applications that it supports. Several applications for this system exist; two are analyzed here.

First, a homeowner's energy costs are a significant part of his/her utilities. If a platform and tracking system could be constructed to help save energy, his/her utility costs will be reduced which saves money. One study [58] shows that by adding some simple sensors and controlling the thermostat, 28 percent of heating and cooling energy can be saved. This amounts to roughly \$15 per month for a typical house [59] and \$180 per year. With additional location-aware algorithms, the

savings could be as high as 40 percent or \$240 per year. If the system costs \$400, \$40 per doorway and 10 doorways in an average house, the return on investment occurs after about two years.

Second, home medical monitoring is an up-and-coming application domain where technology is placed within the home to monitor a patient. Often elderly people wish to maintain the independence and privacy of their own homes and without home monitoring are forced into assisted living facilities. It would be difficult to put a price on the ability to continue living independently. Continuous monitoring of a patient's behavior over a long period of time may help to diagnose diseases that are otherwise difficult to detect. The amount of moving a person does or how many times he/she eats might function as an indicator to depression or other diseases. Both scenarios illustrate applications where knowing the identity and location of individuals allow a system to better monitor their actions and behaviors. The \$300-400 cost of the tracking hardware is significantly below most existing medical care. The average cost of an assisted living facility in Virginia is over \$40,000 per year.

## 3.2 Hardware Design and Operation

The doorway sensors use ultrasonic range finding sensors that are mounted to the top of a doorway and aimed down toward the floor. When the doorway is empty, the sensors report the distance to the floor  $d_f$ . This baseline value indicates the height of the doorway, so in principle the sensors can be installed on doorways of various heights without requiring manual calibration. In practice, however, the baseline measurement is not always of the floor and we manually verify the distance  $d_f$ .

To detect a person's height, the ranging module tries to measure the distance to the top of the head  $d_p$  as a person walks through the doorway. The height of the person  $h$  can be calculated as the difference in these two values:  $h = d_f - d_p$ . The height estimate is most accurate if taken when the person is directly under the door frame. Typically, this corresponds to the tallest height observed, since the  $d_p$  becomes larger as the person moves away from the door frame.

To detect a person's walking direction, the ranging module is angled into one room more than the other, as illustrated in Figure 3.4. This tilt causes an asymmetry in the sensing region, which changes the shape of the curve of  $d_p$  values as the person walks through the door. If the person walks in the same direction as the tilt, the system will detect the tallest heights first, followed by shorter heights. If the person walks opposite the direction of the tilt, the system will detect shorter height followed by taller heights.

The final doorway sensor design consists of five integrated components: (1) two to four Parallax PING ultrasonic range finders, (2) passive infrared sensors facing to each side to sense motion activity in each of the two rooms adjacent to the doorway, (3) magnetic reed sensors to sense whether the door is open or closed, (4) a custom-designed power module to supply regulated 3V and 5V power to the devices, and (5) a Synapse Wireless SnapPY RF100 module, which consists of a Python programmable Freescale processor and an 802.15.4 radio. The components are placed into a molded plastic enclosure designed to mound discretely behind the door jamb. An image of the doorway sensor is shown in Figure 3.2

### 3.3 Challenges

The main challenge for doorway sensor design is to operate under a wide range of doorway sizes, person sizes, and walking speeds while still achieving high height and direction accuracy, few false detections, and few missed detections. Doorway sensors must provide 1 *cm* resolution for people with heights ranging from 151 *cm* (5 ft) to 189 *cm* (6 ft, 2 in), at walking speeds of up to 3 m/sec, in doorways that range from 90-300 *cm* (3-10 ft) wide and 213-275 *cm* (7-9 ft) tall. In order to detect a person, the sensing region must cover the entire width of the doorway and the sensor must sample quickly enough to ensure at least one reading to the top of a person's head when the person is directly under the doorway. The sensing region must also extend outside the door frame in order to detect the direction from which the user enters or exits the doorway. At the same time, the sensing region must not extend too far from the door frame to avoid erroneous measurements such as a person walking by the doorway but not through it.

Because ultrasound sensors have a conical beam angle, it is most challenging to achieve adequate doorway coverage when a tall person walks through a short doorway: the sensing region is smaller at close range, so the person is more likely to walk to the side of the sensing region, or to walk through the sensing region too quickly to be measured (Figure 3.6). The Parallax PING ultrasonic ranging modules used in the hardware design have a 40 degree beam angle, a minimum range of 2 *cm* (0.79 in), and a maximum range of 300 *cm* (9 ft, 10 in). When the tallest person supported walks through the shortest doorway, the gap between the head and doorway sensor is only 24 *cm* (10 in). At that distance, the 40 degree cone width of the PING produces a sensing diameter of 17 *cm* (6.7 in). At a speed of 3 m/sec, a head that is 15cm (6 in) in diameter will pass through this sensing region in about 100 msec, so the ranging module must take measurements at 10 *Hz* in order to ensure




Figure 3.6: Coverage of a pair of ultrasonic PING sensors when placed on a short doorway with a tall person.

at least one reading before the person passes out of the detection region, and a sampling frequency of 30-40  $Hz$  is necessary to produce an accurate and high confidence detection event. The PING module samples at a nominal 50  $Hz$ , but the doorway sensors only trigger one module at a time to reduce multi-path interference. Therefore, a person walking under only one sensor may only be measured at 25  $Hz$ , producing only 2-3 readings before the person exists the sensing region.

At short range, the PING's sensing diameter of 17  $cm$  (6.7 in) is small compared to a typical doorway width of 90  $cm$  (3 ft), and so multiple sensors must be used per doorway. This increases the potential for multi-path interference between modules, and also adds to the financial cost and energy consumption of the doorway sensors. However, the head itself has a radius of about 7  $cm$  (3 in), and so a person can be detected even if the center of the head is slightly outside the PING's sensing diameter. Therefore, the total detection diameter for a single sensor is approximately 30  $cm$  (12 in), and two sensors would adequately cover 60  $cm$  (24 in). Furthermore, the head is highly unlikely to pass through the outermost 15  $cm$  (6 in) on either edge of the doorway because it would nearly hit the side of the doorway due to its 7  $cm$  (3 in) radius. Therefore, a doorway width of 90  $cm$  (3 ft) can be adequately instrumented with only tow sensors, even in the shortest of doorways with the tallest of people.

In high door frames, the detection region becomes larger so even fewer sensors are required for full coverage of a doorway. However, a different set of problems are encountered. When the shortest person supported walks through the highest doorway, the gap between the head and the

doorway sensor is 124 *cm* (4 ft, 1 in), which translates to a detection region with diameter of 83 *cm* (6 ft, 11 in). This large detection region extends outside the door frame and into the neighboring rooms, thereby increasing the chance of false detections. In very wide doorways, on the other hand, sampling rate is limited because only one ranging module is used at a time. One doorway was 300 *cm* wide by 275 *cm* high and required 4 range finders to cover the full span. This reduced the nominal sampling rate from 50 *Hz* to 12.5 *Hz* per sensors. A short and wide doorway is the worst case for sampling frequency: a doorway that is 300 *cm* wide and 213 *cm* high would require 9 ranging modules, reducing the nominal sampling frequency to 5 *Hz* per sensor. At 3m/sec, a person will move 60 *cm* between measurements and could easily pass through the ranging module's 17 *cm* sensing regions without being detected. Thus, the doorway sensor design cannot support doorways that are both wide and short.

### 3.4 Early Prototypes and Lessons Learned

Because of a discrete design space, multi-dimensional optimization is not possible. Instead an iterative design methodology was utilized to design, implement, deploy and test different hardware platforms. Each successive platform combines lessons learned from prior attempts to produce a better system for capturing people as they walk through doorways. This section describes each of these designs in detail.

#### 3.4.1 Version 1.0

Version 1.0 (Figure 3.7) is a commercial off-the-shelf (COTS) product called Go!Motion and was deployed in this author's home, and the data was analyzed by other researchers at the University of Virginia to identify people based on height [22]. This dissertation utilizes this work as a starting point. Go!Motion sensors are typically used in high school physics labs to measure the distance of an object in motion. It was purchased and deployed for a limited amount of time to quickly test the feasibility of an ultrasound based approach.

Go!Motion works well for certain applications. It consists of a computer for control and recording of data along with the ultrasonic sensor which is controlled and powered by USB cables connecting each device to the computer. The software is run on a computer and causes all the sensors to be triggered simultaneously and captures the response time that it converts into distance. Propri-





Figure 3.7: A Go!Motion sensor that was deployed in one of the test houses.

etary software is necessary to record this distance information which limits the use of alternative controllers.

For several reasons, the Go!Motion sensor is not well designed for long term deployment. First, it consumes a lot of power; at least  $200\text{ mW}$  and often more. Second, the sensor produces an audible clicking sound when transmitting a pulse. While this may be acceptable in a lab setting, it is disruptive in a home environment. Third, the control system is based on a USB connection to a computer. In the short two week deployment of just seven devices, hundreds of feet of USB cable were run throughout the house, including across typical walking paths of the occupants. Fourth, the cone width (20 degrees) severely limits the range of visibility when placing a single sensor above each doorway, causing heights to be missed frequently if people did not walk directly under the sensor. Fifth, the  $20\text{ Hz}$  maximum sample is mathematically enough to capture a person but practical rates should be higher for more confidence and the ability to determine more features such as direction

from the signals. Lastly, the aesthetics of the design are less than desirable. The Go!Motion sensors are highly visible bulky black boxes. Because the deployment was a short one in a researcher's home, the boxes were taped to the walls which is something that probably would not have been allowed otherwise. These limitations make this device impractical and unable to accurately capture the human motion under the doorways.

### 3.4.2 Version 2.0

Version 2.0 (Figure 3.8) moves away from simple systems in favor of a custom design that incorporates different ultrasonic hardware and a wireless sensor node for control and data collection. This particular ultrasound sensor is typically used for measuring the level within various industrial containers such as silos or fluid tanks where manual measurements are challenging.



Figure 3.8: A Maxbotix and TelosB height sensor deployed inside a doorway.

Two different hardware devices compose the version 2.0 hardware platform. First, a typical wireless sensor network node, the TelosB [60] reads distance measurements produced by the ultrasound device and transmits the readings wirelessly to a central laptop. It is connected to the Maxbotix (XL-MaxSonar-EZ0) [61] piezoelectric ultrasound transducer to measure distance. Capable of a 10  $Hz$

sampling rate with a 90 degree cone width and measuring distances up to 25 feet due to the high receiver gain and large signal energy, this platform is better suited than the Go!Motion platform to measure the desired features as people walk through doorways; however, some challenges do exist.

A number of problems appeared after deploying the Maxbotix sensor system. First, the devices produced too much acoustic energy which results in interference between doorways. Large amplifier gain and cone width on the transmitter result in signals reflected off the floor and into another doorway. Clock drift between separate TelosB nodes that were responsible for taking measurements caused the interference patterns to change over time as the phase of sampling changed. There are two possible solutions for this destructive interference: either use a precise software time synchronization protocol, or the chosen solution which was to run a 250 foot wire and perform synchronization in hardware.

Second, the sample rate was software-limited by the manufacturer to 10 *Hz* for this particular device due to its long range which implies a long RTT. Barely above the minimum sample rate needed to detect people in doorways, this typically resulted in only two samples recorded for each event which was well below the necessary points to reliably compute both height and direction. These two issues resulted in loss of events either from the abundance of noise or the lack of data points.

Several lessons were learned from the version 2.0 hardware. 10 *Hz* is not a sufficient sample rate to isolate a height measurement of a moving person. Clock drift and ultrasound sensors with too much transmission power result in cross-doorway interference that can render large blocks of data unusable. An analysis shows that for certain doorway pairs, up to 16 hours of data per day was rendered unusable because of this error. Unsightly wires (250+ feet) and exposed circuit boards are aesthetically challenging and only accepted for a limited time by house occupants. Such devices were only tolerated for approximately six months and would have been removed had the house occupants not been part of the research group. These lessons are taken into account and form the basis for the final version of the hardware.

Version 2.0 provided more data (6 months) but identified some of the more subtle challenges associated with using ultrasound transducers to measure a person walking through a doorway. The development time was slower than that of the Go!Motion devices; however, deploying and debugging the hardware was much more difficult because of the custom solution. A number of drawbacks to the design appeared and will be addressed in the final version of the hardware.

### 3.4.3 Version 3.0



Figure 3.9: A Version 3.0 sensor deployed in a normal doorway.



Figure 3.10: A Version 3.0 sensor deployed in an atypical doorway.

Version 3.0 (Figures 3.9, 3.10 and 3.11) is the final hardware design, as described in Section 3.2 and incorporates the lessons learned from previous approaches. The platform design integrates several off-the-shelf components. First, a Synapse Wireless SnapPY (RF100 [62]) is the new choice for controlling the sensor platform. It runs Python which makes developing code much faster than the TinyOS [63] NesC based systems. It supports over-the-air (OTA) programming allowing application code to be remotely changed and tested without needing to physically connect to each device. These two factors combine to allow for the rapid development of algorithms and to make integration of new sensors easy. Second, multiple PING [64] ultrasonic modules can be attached to a single RF100

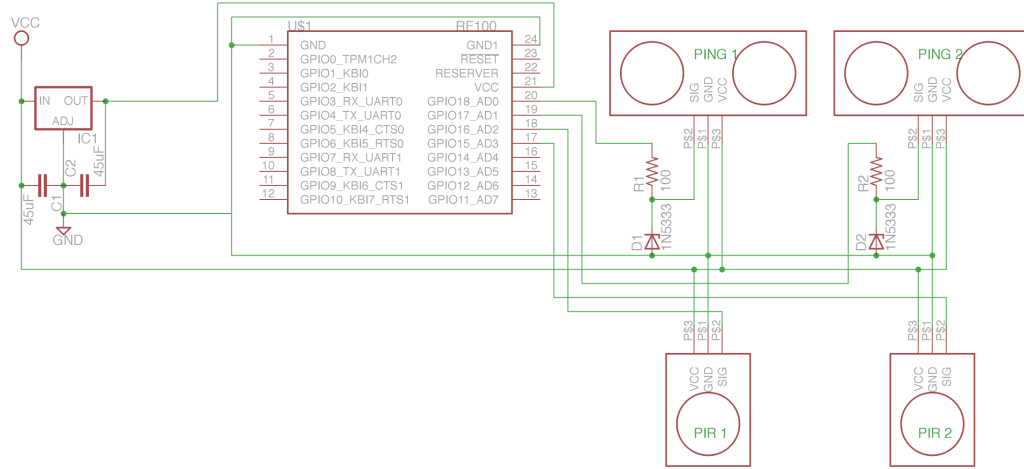


Figure 3.11: Height sensor version 3.0 schematic.

device. The PING sensors also have several challenging characteristics for human experiments: (1) their measurement range is sufficient for people walking underneath sensors placed above the door and (2) the sample rate is not limited by software which allows it to sample as quickly as possible; thus it has a dynamic sample rate that increases the closer an object is to the sensor. However, the maximum sample rate must be shared between all PING sensors on each doorway because they sample in a sequential manner to avoid interference problems. The energy used is not sufficient to cause intra-door interference. Multiple PINGS with narrower cone widths reduce the effects of geometric error in the measurements.

Third, two other types of sensors are included for future work: (1) passive infrared (PIR [65]) sensors to measure motion placed facing out from the doorway and (2) magnetic reed sensors to indicate whether a door is open or closed. Fourth, a custom-designed power module allows us to adjust the voltage driving the PING transmitter, increasing its range to desired distances when mounting on non-standard doorway heights.

The components are placed into a molded plastic enclosure to be more aesthetically appealing so that the deployments will be accepted for longer periods of time. Custom-designed circuitry allows the RF100 hardware platform to interface with all the sensors that will be placed on each doorway. Each component plays an important role in how the final product will be used and what data will be collected which affects the reliability of the device and system.

The deployment of the final version taught more lessons than both of the prior systems. First, design principles which *disappear into the woodwork* received praise by house owners and occupants who were familiar with the other systems. Second, multiple ultrasound devices per doorway helped with geometric error and missing events but caused measurement noise when the sensors interfered

with each other. Third, taller than normal doorways resulted in a large drop in the sampling rates because the round-trip times for sound pulses were much longer. This was not a problem when there was only a single sensor per doorway; however, one doorway had four PING devices and only achieved a net sample rate of about 13 *Hz*. Fourth, the hand-soldered circuit boards were not robust components and could fail at deployment time resulting in increased debugging time. In future iterations, it would be worth the effort and cost to have circuit boards printed and assembled by machines. These lessons learned should be applied to future work when designing new platforms for home deployments.

Version 3.0 is the final design and was implemented and installed in over 40 doorways in four different houses and achieved an accuracy of 0.5 *cm*. Each house had a variety of challenges and doorway configurations necessitating the need for a custom fit device for each. It is difficult to anticipate the challenges discovered at each stage of development to arrive more quickly at the final solution. Initially, this project was supposed to be solved in two months, but due to the many complexities of monitoring people in an indoor environment, it developed into a much longer project taking over a year of experiments and deployments.

## 3.5 Alternative Ranging Technologies Considered

Several alternative ranging technologies to the PING ultrasonic range finder were considered but found to be inadequate for the smart home application. With further development, many of these alternative technologies could be equal to or better than the ultrasonic solution used, but not in the forms currently available. This section summarizes the findings and explains the shortcomings of each technology, in comparison to the final solution.

### 3.5.1 Laser Range Finding

Laser range finding solutions fall into two distinct classes. The first approach utilizes single point sensors and measures distance in a very narrow area by recording different frequency shifts of the returned signal. Second, light detection and ranging (LIDAR) systems measure distance based on Time-of-Flight (TOF) and employ additional hardware allowing them to measure whole environments.



### Multiple Frequency Phase-Shift

Many off-the-shelf technologies utilize the multiple frequency phase-shift (MFPS) technique to determine how far a target is from the source. The laser is modulated on multiple frequencies and the reflected signals are used to compute the distance to the object. For example, a golfer uses a hand-held device to measure the distance between the ball and flag or a builder uses it to quickly measure the dimensions of a room without relying on a tape measure. These devices rely on the measurements occurring between specific points.

Laser-based devices (Figure 3.12) have several desired properties. They are single point measurement devices easily interfaced with low power controllers. Second, the extremely narrow laser beam limits most interference from outside sources or other measuring devices. These two desirable characteristics serve to eliminate sources of noise.



Figure 3.12: Single point laser system measuring the distance in a building under construction.

Because they are typically used outdoors to make measurements over large distances, creating laser-based systems precise enough for indoor use is difficult and expensive. This is because light traveling at  $3.0 \times 10^8$  meters per second requires precise timing to obtain accurate measurements. A precise point and narrow beam width result in a limited amount of area seen by the signal which makes interference either caused or received by this class of sensors highly unlikely. While the capabilities of TOF laser devices appear to meet most of the hardware design requirements, a number of challenges do exist when utilizing these devices indoors for measurements in the 0 to 100 *cm* range.

First, using low power devices makes the timing accuracy for the speed of light difficult or expensive. Achieving a 1 *cm* accuracy requires a measurement precision of  $50 \times 10^{-12}$  seconds with

the round-trip time for a distance of 15 *cm* being 1 nanosecond. Considering that the clock rate for the controller nodes is typically about 16 *MHz*, these timings are difficult to achieve.

Second, narrow beam width sensors only measure the part of the body directly underneath the sensor. For example, a 1 *cm* offset in a person's head could cause a narrow beam sensor to measure the shoulder which results in an error of about 30 *cm*. Because the design is based on measuring height and not just any part of the body additional sensors are required and drive up the cost and energy requirements of the system. These two challenges make laser MFPS systems expensive to implement and the necessary timing circuits are not low power enough.

MFPS laser systems work great where distance is needed but centimeter level precision is not required. These systems are desirable because their design limits the amount of interference from the environment and other devices. With appropriate timing circuits, these systems can be highly accurate, but these circuits are expensive and power demanding.

### **Light Detection and Ranging**

A variety of robotic and industrial environments such as Google's self-driving car, utilize Light Detection and Ranging (LIDAR) (Figure 3.13). LIDAR provides information to keep the car from hitting anything or driving off the road as the primary navigation system. Also used for three-dimensional modeling and scanning systems, LIDAR systems accurately measure building architecture and generate a high-precision three-dimensional model. There are two useful characteristics of this class of system: three-dimensional point clouds and accurate measurements. Several drawbacks exist because of the high accuracy and detailed information including cost, power, and data rate which makes them impractical for monitoring people moving through doorways.

LIDAR systems scan an environment and construct a three-dimensional point cloud by measuring the distances to each point seen by the device. There are a few commercial LIDAR systems, including one which produces up to 1.3 million points per second and utilizes a 100MBit/s communication interface. LIDAR systems provide extremely rich information about their environments.

LIDAR accuracy is determined in a similar way to single-point TOF systems. The most expensive device on the market that can sample above the 6 *Hz* minimum sample rate only has a 2 *cm* distance accuracy, but because of the intended purpose, has a minimum sample distance that is larger than the maximum necessary distance. Any hardware chosen should far exceed the minimum sample rate but this device is limited to a maximum of 15 *Hz*. This is close to the design goals for precision but falls short with the minimum sampling distance and rate.






Figure 3.13: The Light Detection and Ranging Sensor (LIDAR) that is used in Google’s self-driving car.

Indoor person measurement applications present three challenges for LIDAR systems. First, costing anywhere from \$2,000 to \$75,000 per sensor, they are too expensive to install on every doorway. Second, these devices consume a lot of power, up to 15 W, due to the mechanical and laser components. Finally, the high data rate produces too much data for anything but a desktop computer to handle.

Like laser MFPS, LIDAR systems are extremely useful in outdoor or static environments, but are inappropriate for indoor use when monitoring people moving through doorways. They produce significant data that needs to be processed with a computer and consumed more power in a single device than the anticipated power budget for the whole house. Additionally, their cost is prohibitive for installation on each doorway.

### 3.5.2 Infrared Range Finding

Infrared (IR) range finding technology is similar to laser TOF techniques; however, the angle of reflection is measured. A cheap sensor (\$10) utilized by robotic vehicles can detect how far they are from obstacles. While providing similar benefits to laser-based systems, infrared sensors cost less and consume less power. They also present different challenges.

Distance can be measured by determining the *Angle of Reflection* of a transmitted signal back to a receiver (Figure 3.14) and then triangulating it. The angle of reflection is the angular difference between the axis of sensing and the received signal back at the device. The closer an object is to the sensor, the greater the reflection angle. An easy way to implement this measurement is to utilize a IR transmitter and a lens plus a linear optical sensor. Due to the need to measure reflect angle, these sensors utilize a narrow beam width transmitter to prevent any stray light from causing problems.




Figure 3.14: The distance from the sensor to an object is determined by the angle an infrared light pulse is received.

The precision of measurement is determined by the reflection angle. This results in a decrease in accuracy the further away from the sensor an object is located. If an object is too close, it will either appear to be further away or not be measured. This optical sensor can be built to various specifications, but they commonly only measure in the 25-100 *cm* range with a precision of 1-2 *cm*. This is exactly what the design goals have specified.

Several new challenges arise when utilizing angle of reflection as the basis for measurement. First, accuracy falls off at large distances resulting in the precision to drop below 10 *cm*. Second, a variety of surfaces create problems for infrared reflection. Human hair tends to absorb the infrared signal causing the maximum sensing range to drop. Finally, multiple sensors per doorway are required due to the limited beam width. These difficulties are identical to those of single-point laser systems, making it difficult to use this type of sensor in a home with uncontrolled lighting conditions and a variety of distances to measure that may or may not fall into its sensing range.

Infrared range finding is highly effective at precisely measuring the distance to obstacles in the robotic domain; however, measuring human heads with hair is more difficult as they move through doorways. While these infrared sensors are cheap, they are restricted by narrow beam width and

consume approximately 100 *mW* of power per device. 8-10 devices are necessary to cover a doorway resulting in at least 1 *W* of power per doorway.

### 3.5.3 Infrared Point Cloud

Infrared point cloud systems generate a three-dimensional model of an environment. The most famous device is the Kinect<sup>TM</sup> (Figure 3.15) sensor which is used to play video games with the human body and utilizes no physical controller on the XBox360<sup>TM</sup>. The infrared point cloud systems are similar to the LIDAR systems because they both scan the environment and produce a multi-dimensional array of measurements. However, each device acquires the measurements in a different manner.

The Kinect<sup>TM</sup> works by using an infrared projector to shine a grid of points into an environment which are then recorded and the deformation of points measured. Distance is computed from the difference between expected and measured point locations. It operates with a video camera at up to 30 frames per second which allows for the capture of quick movements. Additionally, a color camera is used to help with feature computation and image segmentation for use by the XBox360<sup>TM</sup>. These characteristics allow this class of devices to capture a lot of information about the environment.



Figure 3.15: An infrared transmitter transmits a grid of infrared points into the environment while a camera records these points and determines the distance to each.

Some challenges do exist. Measurement accuracy is a function of the resolving power of the infrared camera, similar to the precision of the angle of reflectance devices. Specifications of the Kinect<sup>TM</sup> device indicate the best obtainable accuracy is 2 *cm* at 30 *Hz* with a sensing range between 120 *cm* and 350 *cm*. The accuracy of 2 *cm* does not quite meet the design goals and the minimum sensing distance falls short.

Infrared point cloud systems face a number of challenges for indoor person monitoring. First, the minimum sensing range of the Kinect<sup>TM</sup> is not tuned to the proper range for the application space. Second, these devices are high-power (15 W) and incapable of being powered by batteries for any significant period of time. Third, a lot of processing power is required to handle the data rate and compute useful features from the data. These challenges make implementing the system difficult with the limited power budgets and other resource constraints of a micro-controller.

Infrared point cloud techniques provide useful information about people walking not only through a doorway, but through any indoor environment in a cheap yet data rich setup. However, they also have a number of drawbacks which include power consumption and privacy invasion from the video cameras when attempting to use this device in houses continuously over a long period of time.

### 3.5.4 Electrostatic Ultrasound

Electrostatic ultrasound sensors measure distance by transmitting an acoustic pulse and recording its time-of-flight (TOF,) which is most similar to the final solution. An example of an electrostatic ultrasound device, the Go!Motion [66] sensor (Figure 3.16) is capable of meeting most of the design goals though there are some drawbacks.



Figure 3.16: The Go!Motion electrostatic ultrasound device works by transmitting an acoustic pulse and measuring how long it takes the pulse to return.

Ultrasound sensors work by transmitting an acoustic pulse and measuring the amount of time until the reflection returns. The hardware typically consists of either a transmitter and receiver pair or a transducer which performs both jobs. Electrostatic devices employ a high-voltage ( $\pm 200$  volts) applied across a diaphragm to generate an ultrasonic pulse. Their designs usually translate into wider beam widths of 20 degrees or more when compared to other optical-based sensors. The electrostatic devices produce cleaner signals when compared to other classes of ultrasonic measuring devices due to the amount of energy put into each transmission by the high voltage.

A large variety of challenges exist for this class of devices. Even with a 20 degree cone width, multiple sensors per doorway are required to avoid measuring the wrong parts of a body; therefore, fewer are required than with the optical systems and the total power consumption is reduced. Electrostatic devices produce a characteristic clicking sound when transmitting a pulse. While this click is acceptable in a laboratory setting, it is disruptive to the home environment. Reflected noise, ranging noise, and interference from other devices create additional issues when using ultrasound instead of light. These are addressed in detail later in this dissertation. Even with all the problems and challenges associated with electrostatic ultrasound devices, they cover more of the design goals set forth than any other approach. Electrostatic ultrasound devices are not suitable for indoor home deployments because their power consumption is too high to feasibly run them off of batteries and their audible clicking is not tolerable by the people occupying the house for significant periods of time.

### 3.5.5 Piezoelectric Ultrasound

Piezoelectric ultrasound sensors measure distance identically to the electrostatic sensors with TOF. An example of a piezoelectric ultrasound device, the PING [64] sensor (Figure 3.17) is capable of meeting most of the design goals though there are some drawbacks.



Figure 3.17: The PING piezoelectric ultrasound device works by transmitting an acoustic pulse and measuring how long it takes the pulse to return.

Piezoelectric devices employ a low-voltage ( $\pm 10$  volts) applied across a piezoelectric transducer to generate an ultrasonic pulse. Their designs usually translate into wider beam widths of 90 degrees or more when compared to other optical or electrostatic-based sensors. The signals produced are not as clean because of the lower voltage; however, they do not incur the same drawbacks as electrostatic devices. They do not produce any audible noise which does not annoy the people living in a house. However, it was discovered that these devices are not cat friendly because of their extended hearing range and can not be used in houses where they live. The other benefit is their power consumption is significantly lower than the electrostatic-based devices.

## 3.6 Evaluation

There are three different ways to evaluate the hardware designs. First, the alternatives are examined and evaluated in five categories: aesthetics, accuracy, cost, power, and range. Second, a power analysis is performed on the three versions of hardware deployed in test houses. Finally, a breakdown of the costs for Version 3.0 is provided along with an evaluation of the return on investment that could be expected for some scenarios.

### 3.6.1 Platform evaluation

The evaluation of our hardware solution takes on a qualitative instead of a quantitative approach. The individual platforms are qualitatively evaluated over five categories: aesthetics, accuracy, cost, power, and range.

| Platform             | Rate             | Aesthetics | Accuracy       | Cost     | Power (watts)     | Range             |
|----------------------|------------------|------------|----------------|----------|-------------------|-------------------|
| Kinect <sup>TM</sup> | 30 <i>Hz</i>     | Low        | 1-2 <i>cm</i>  | \$150    | 15 <i>W</i>       | 120-335 <i>cm</i> |
| Sharp IR Laser       | 100+ <i>Hz</i>   | High       | 1-10 <i>cm</i> | \$140    | 2.1 <i>W</i>      | 20-150 <i>cm</i>  |
| LIDAR                | 100+ <i>Hz</i>   | Low        | 2 <i>cm</i>    | \$2,000+ | 60+ <i>W</i>      | 3-120 <i>m</i>    |
| Go!Motion            | 20 <i>Hz</i>     | Low        | 1-2 <i>cm</i>  | \$100    | 250 <i>mW</i>     | 15-600 <i>cm</i>  |
| Maxbotix             | 10 <i>Hz</i>     | Moderate   | 1-2 <i>cm</i>  | \$120    | 150 <i>mW</i>     | 20-700 <i>cm</i>  |
| PING                 | 50-500 <i>Hz</i> | High       | 1 <i>cm</i>    | \$95     | 150-400 <i>mW</i> | 2-300 <i>cm</i>   |

Table 3.1: Evaluation of a spectrum of height sensing solutions along six variables (sample rate, aesthetic appeal, accuracy, cost, power, and range). Alternative solutions have at least one metric that is not as good as the final solution, designated as PING.

Table 3.1 summarizes the hardware designs explored to evaluate alternative approaches in six categories. Sampling rate is evaluated based on how quickly the sensors can acquire the next reading.

Aesthetic appeal is evaluated in Column 2 and is a subjective evaluation of how well each design will fit into a house environment. The Kinect<sup>TM</sup> sensor received the low rating but not because of its design. It incorporates video cameras into its sensor suite which none of the houses were willing to accept. The Go!Motion devices produce an audible click when taking measurements which was an unknown problem at the time of purchase. In a single house trial that lasted two weeks, they were limited to running only during the day and had to be turned off at night. LIDAR-based systems were not utilized because their form-factor is typically too large to place in a doorway. Laser-based systems received high marks in the visual category because they are small and easily hidden in enclosures like ultrasound devices. The three ultrasonic approaches progressed from low to high aesthetic appeal through the various iterations to *disappear into the woodwork*.

Sensor accuracy measures the precision of the sensor and is summarized in Column 3. Some of these values were determined from specification sheets as the sensors were not purchased and tested directly. The laser systems in particular have poor measurement accuracy which varies based on the distance. The rest of the platforms perform well with respect to accuracy and meet the design goals.

Column 4 shows cost which is an important factor, especially for energy conservation applications where the return on investment is important. Some of the sensors are extremely expensive and can cost up to \$75,000 in the case of LIDAR while others cost about \$100. Five of the six cost between \$100 and \$150 making an eight doorway house cost about \$1,000.

Power consumption is a consideration for the future of this platform when batteries or solar harvesting will be the only power source, making high power devices impractical. Kinect<sup>TM</sup>, laser, and LIDAR consumed more than 1 *W* of power, with the LIDAR systems breaking 60 *W*. These power requirements are too high for batteries that can be placed in an enclosure in the doorway. The three ultrasonic solutions consume between 150 and 400 *mW* which would allow between 6 and 16 hours of operation on batteries. This means that the batteries of a sensor running continuously must be changed more than once per day; however, assuming 1 percent utilization which is the approximate amount of time a person is under a door frame each day, the batteries need only be changed every 42 days.

Minimum and maximum sensing range is a critical metric because it must measure in the correct range of a doorway. Neither the Kinect<sup>TM</sup> or LIDAR systems meet the minimum requirements for distance. They have a minimum distance starting at 1.2 meters resulting in the inability to measure the short distances to the top of a person's head. The laser-, Go!Motion-, and Maxbotix-based systems all have minimum distances that work for most doorways; however, even these minimums

may sometimes be too much if an occupant is tall enough to bring the distance between the sensor and head to about 10 *cm*.

All six approaches have desirable properties. For the prototype, a solution was chosen that offered a good balance between competing goals.

### 3.6.2 Ultrasound Platform Power Analysis

The power consumption of various ultrasound hardware solutions var widely based on specific abilities. Each of the three versions deployed and tested is evaluated in three different running configurations (Table 3.2). First, Version 1.0 is composed of only the Go!Motion sensor system and consumes 255 *mW* of power when running. This particular device is unable to be duty-cycled. If the sensor could be utilized in an optimal manner when a person is walking under the doorway, the estimated power consumption will be 3 *mW* because the sensor only needs to be on approximately 1 percent of the day. These two power measurements do not account for the power usage of the computer that is necessary to drive the devices.

| Platform    | Sensor    | Power           | Duty-cycle     | Ideal          |
|-------------|-----------|-----------------|----------------|----------------|
| Version 1.0 | Go!Motion | 255 <i>mW</i>   | N/A            | 3 <i>mW</i>    |
| Version 2.0 | TelosB    | 82 <i>mW</i>    | 10 <i>mW</i>   | 4 <i>mW</i>    |
| Version 2.0 | Maxbotix  | 12.5 <i>mW</i>  | 10 <i>mW</i>   | 8 <i>mW</i>    |
| Version 2.0 | Total     | 94.5 <i>mW</i>  | 20 <i>mW</i>   | 12 <i>mW</i>   |
| Version 3.0 | RF100     | 100 <i>mW</i>   | 50 <i>mW</i>   | 10 <i>mW</i>   |
| Version 3.0 | PING      | 150 <i>mW</i>   | 20 <i>mW</i>   | 10 <i>mW</i>   |
| Version 3.0 | PIR       | 0.5 <i>mW</i>   | 0.5 <i>mW</i>  | 0.5 <i>mW</i>  |
| Version 3.0 | Latch     | 0 <i>mW</i>     | 0 <i>mW</i>    | 0 <i>mW</i>    |
| Version 3.0 | Total     | 250.5 <i>mW</i> | 70.5 <i>mW</i> | 20.5 <i>mW</i> |

Table 3.2: Each of the three iterations of sensor design is broken down by component power for three different sensing scenarios: always on (Power), duty-cycled by periodically waking up and sampling, and optimal, where the sensor is only on when an event is occurring.

Second, Version 2.0 is built on the TelosB and Maxbotix platform and consumes a total of 94.5 *mW*. The bulk of this power consumption goes into the TelosB device, but techniques such as turning off the radio or placing the micro-controller into a sleep state could be used to reduce power. When this is done, the estimated power consumption drops to 20 *mW* by waking up periodically and taking a measurement. If this device were able to only wake up and sample when a person was under the doorway, its power consumption would drop further to 12 *mW* because the node will only be in an active state exactly when a person is under the door way. This is defined as the *ideal* power.



However this particular device cannot approach the optimal consumption rate without additional hardware.

Finally, Version 3.0 is built on the RF100 node and the PING ultrasound transducers. Its initial power consumption is higher than Version 2.0 at 250 *mW*, mainly due to the multiple PINGs required for each doorway and the PING device consumes more power than the Maxbotix device; however, it is capable of sampling at a much higher rate. By duty-cycling this hardware, the estimated power consumption drops to 70mW. Most of the savings occur on the PING devices because they are not continuously transmitting. Since this device also contains PIR and Latch sensors to better manage power consumption, it should approach an ideal power level of 20 *mW* where a sensor is only active when a person is under the doorway. Although this version is not the most power efficient platform, it provides the necessary sample rates and additional sensors to better measure people.

### 3.6.3 Cost analysis

| Part Name              | Unit Cost | Cost (10000) | Req. | Prototype | House  | Production |
|------------------------|-----------|--------------|------|-----------|--------|------------|
| AC to USB converter    | 2.00      | 1.00         | 1    | 2.00      | 20.00  | 10.00      |
| Male USB plug          | 0.4052    | 0.20250      | 1    | 0.41      | 4.10   | 2.03       |
| 3.3V voltage regulator | 1.60      | 0.80000      | 1    | 1.60      | 16.00  | 8.00       |
| 100 ohm resistor       | 0.0271    | 0.00992      | 2    | 0.05      | 0.50   | 0.20       |
| 3.3V Zener Diode       | 0.078     | 0.02535      | 2    | 0.16      | 1.16   | 5.07       |
| 45uF Capacitor         | 0.375     | 0.19500      | 2    | 0.75      | 7.50   | 3.90       |
| 2mm header pins        | 1.16      | 0.941        | 2    | 2.32      | 23.20  | 18.82      |
| Circuit board)         | 0.55      | 0.25         | 1    | 0.55      | 5.50   | 2.50       |
| Enclosure              | 1.56      | 0.80         | 3    | 4.68      | 46.80  | 24.00      |
| 4 strand wire          | 0.07      | 0.035        | 25   | 1.75      | 17.50  | 0.88       |
| SNAPPY                 | 24.00     | 20.00        | 1    | 24.00     | 240.00 | 200.00     |
| PING sensor            | 20.00     | 15.00        | 2    | 40.00     | 400.00 | 300.00     |
| PIR sensor             | 6.50      | 5.00         | 2    | 13.00     | 130.00 | 100.00     |
| Latch sensor           | 4.10      | 2.50         | 1    | 4.10      | 41.00  | 25.00      |
| Total Cost             |           |              |      | 95.37     | 953.70 | 700.40     |

Table 3.3: All the components that make up the final hardware platform. Each component is priced for individual purchase and the cost point for quantities of 10,000. Then the cost of constructing an individual prototype is compared to the cost to outfit a typical house with 10 doorways, followed by the expected production cost for deploying in the same house.

Cost has been one of the driving factors when constructing this hardware platform because energy efficiency applications that rely on a tracking system focus on the return on investment. The bill of materials (Table 3.3) prices the components used for both a prototype system and for a large scale production run of 10,000 units. The cost of a single prototype sensor for a typical doorway is

\$95.37. This is comparable to other solutions in this class of sensor systems. The entire system costs \$935.70 when deployed in a typical house containing ten doorways. The expected production cost for producing this device in quantities of 10,000 and deploying in a ten doorway house is \$700.40. A new custom design that does not use pre-integrated parts could cost as little as \$300.00 per house.

This price point of less than \$1,000 allows the system to be easily purchased by a homeowner. For example, if a location-aware heating and cooling system could save \$200 per year, the return on investment is 3.5 years. Additionally, this system can be used simultaneously for other purposes such as medical monitoring and activity recognition which further adds value.

### 3.7 Conclusion

A variety of devices were explored with three different ultrasonic devices deployed into houses for testing. This chapter presents an analysis of the hardware designs along with the lessons learned to illustrate and motivate the next design. Piezoelectric ultrasound was chosen from a large class of possibilities based on a qualitative evaluation in term of aesthetic appeal, accuracy, cost, power, range, and sample rate.

Laser MFPS and infrared sensors require too many devices on a doorway to measure the height properly. Additionally, laser systems require extremely precise timing circuits to measure to the desired accuracy. LIDAR and point cloud systems produce too much data to be used by a low power sensor platform. Each requires significant processing to acquire the height information in a useful form. Electrostatic ultrasonic devices are most similar to the final approach, but have characteristics that prevented their use: high-voltage transducers and audible clicking when sampling. Each of the systems has benefits and drawbacks. Piezoelectric ultrasonic devices provide a reasonable balance among the design goals.

There is still much work to be done when measuring and recording human motion in doorways. PIR and latch sensors were included as part of the latest design in the hope that the data will be helpful for future algorithms. The noise problems with ultrasound will not disappear through incremental improvements in the device and something new will need to be devised to measure the motion more precisely and at a higher sample rate.

## Chapter 4

# Signal Processing

Signal processing should detect and record a person walking under the doorway and determine both height and direction.. It addresses a multitude of noise challenges within the data due to physical characteristics of the people monitored and the environment where sensors are located.

Prior attempts [22] at utilizing ultrasonic devices to measure height have done so only in lab or isolated settings. These limited deployments result in easy-to-process data where a simple threshold filter is sufficient to locate where events occur. The hardware design in this dissertation has been in place for 18 months, resulting in over 17 billion readings. The physical environment is continually changing which alters the types of noise seen and makes identifying where events occur difficult. The hardware is designed to measure the height of an adult person at each doorway. An artifact of this solution is a tendency for one of the two sensors to report invalid heights, indicating a person is there but not a valid height. This artifact is difficult to detect and requires a more sophisticated algorithm.

The signal processing algorithms need to solve four challenges to correctly identify events, heights, and directions: environmental interference, environmental change, target interference, and sensor noise. These challenges appear in their own unique ways within the data and are handled in slightly different ways by signal processing algorithms.

Experimental results indicate that the algorithms can identify 96 percent of the events within a particular data stream. Typically 10 samples long, these events are correctly identified within the context of 180,000 samples per hour for each doorway. The algorithms also detect the direction a person is moving during an event with up to an 87 percent accuracy. This chapter presents the challenges associated with ultrasonic data processing in doorways and discusses the specific

algorithms and implementation used to handle these challenges. Additionally, an evaluation of the effectiveness of signal processing is done with a series of empirical experiments. This chapter is published as part of a *Sensys* 2012 paper.

## 4.1 Challenges

There are several challenges associated with processing sensor data to determine events and compute desired features. These challenges fall into four groups. First, *environmental interference* occurs when physical characteristics of the environment where the sensor is mounted affect the measurements. Second, *environmental change* occurs because houses are dynamic spaces where people continually reconfigure the space. Third, *target interference* occurs when noisy values are generated by parts of a person other than the top of his/her head. Finally, *sensor noise* occurs when measurements returned by the hardware are incorrect due to imprecision of the range sensor. These four challenges should be addressed by the signal processing algorithms.

### 4.1.1 Environmental Interference

Environmental interference is caused by objects in the environment other than the target and appears as noise bands in the data. Figure 4.2 identifies three different example sources of noise. First, a reflection from door molding can occur in the 150-200 *cm* band that arrives back at the receiver. Second, a door knob reflection is common when doors are closed and is indicated in the figure as the middle noise band. Finally, floor and base molding typically occur in various forms for most sensors and is shown in the lowest band.

There are two major problems with environmental interference as it relates to noise filtering. First, noise sources vary by doorway, and each doorway in a house will have a different signature. Second, it is difficult to determine the noise sources until after deployment. These two problems show a need for an algorithm that adapts to varying noise bands without any previous knowledge about the physical environment.

Noise bands may change over large temporal periods due to physical environmental changes. Objects such as a box, couch, or toy that happened to be left on the floor may be moved in and out of doorways (Figure 4.1). Over longer periods of time, the sensor mounting may also shift changing the noise produced by the environment. Some of the noise sources occur based on whether a door is open, closed, or partially ajar. For example, Figure 4.1 illustrates three different hazards that contribute to multi-modal noise. A doorknob about waist high easily reflects the signals causing a

noise band in this region. This band changes height depending on how far the door is left open. If the sensor can pick up the counter top to the right of the door, it will also cause noise at about the same height as the door knob. A picture placed on the floor against the door frame can easily be picked up by the sensors, and because it is temporary, its noise band will be unknown. Problems like environmental interference illustrate the need for algorithms that can differentiate environmental change from people walking through doorways.



Figure 4.1: Environmental Hazards

Figure 4.2 shows the front door sensor in one house. It is sandwiched between a screen door and a solid wood main door with only about 15 *cm* of clearance. Up until 11:45, there were two large modes located at 0 *cm* and 180 *cm* that correspond to where the sensor is sampling the floor and door frame. A third mode is created when the door is opened and is most likely a door handle or decorative trim on one of the doors.




Figure 4.2: Multi-modal noise occurs when there are multiple objects at different levels that the sensors pick up. Arrows point toward three large modal bands. Each band contains tens of other modes.

#### 4.1.2 Measurement Error

Motion, hair, and skin change the reflected waveform resulting in measurement error. Furthermore, other environmental factors such as carrying a cellular phone or different parts of the house add acoustic noise and additional reflections. There are many reflected signals including multi-path signals and the PING device needs to choose which one is the line of sight reflection. It does so by matching the signal delay with its expected amplitude. Some surfaces have different reflective properties (size of reflected surface, angle, material, etc.), which cause ambiguity with the matching process. Figure 4.3 illustrates how frequently noise occurs and how many times heights are recorded. Relevant data is identified by the two arrows indicating *Person 1* and *Person 2*. These spikes indicate height readings appropriate for each person. Additionally, the three spikes indicate physical objects located at roughly 0 cm, 20 cm, and 40 cm that are due to moldings on the door and door frame. Noise is indicated by the Gaussian-shaped signals surrounding these objects and generally ranges from 20-40 cm.

Target interference is caused by carried objects such as a box, cellular phone, child, or parts of the body that cause additional reflections such as ears or shoulders. Figure 4.4 illustrates a person




Figure 4.3: Baseline sensor noise is shown on the left hand side of the figure. Given the three large spikes, which are stationary objects and do not move, the variance around them is the noise.

walking through a doorway as he/she carried a cellular phone face up at waist level. This introduces a secondary spike of heights that could confuse a threshold-based filtering algorithm.

People walk through doorways in a variety of different ways using different speeds, postures, and angles. Speed affects measurement error because a person moving faster is less likely to be directly under a sensor when the measurement occurs. If a person slouches or is carrying a child, his/her posture is different and results in a lower height measurement. Finally, the angle a person walks through a doorway could affect how well a sensor detects his/her height because he/she might not be directly one of the sensors.

Geometric offset error causes heights to appear shorter than they actually are and is a function of where a person walks through a doorway. If the head passes directly under the sensor, geometric error is minimized; however, most of the time, the head does not pass directly under the sensor which introduces additional error.

The final source of measurement error comes from multi-path signal reflections that occur when signals reflect off multiple surfaces and multiple sensors pick up the same signal resulting in incorrect height measurements. This error is typically seen in the data streams as heights above the actual height of a person and indicates that a reflected signal was picked up by a second sensor before it received its own reflection. These four challenges illustrate the need for robust signal processing




Figure 4.4: This is an event where a cellular phone was being carried in front of the body while someone walked through a doorway. The cellular phone height is indicated with the circled data points.

algorithms that can adapt to unknown signal errors and isolate true height measurements from the noise

## 4.2 Approach

A doorway sensor with  $k$  range finders generates a  $k$ -valued column vector  $X$ , where each element  $X_i$  at time  $t$  is a continuous stream of  $t$  height values produced by the  $i$ th range finder:  $X_i = x_i^1, x_i^2, x_i^3, \dots, x_i^{t-1}, x_i^t$ . The goal of the signal processing component is to convert this vector  $X$  into a single set of doorway events  $D$  where  $d_j \in D = (t_j, h_j, v_j)$ : a 3-tuple containing the timestamp of a doorway crossing  $t_j$  and the corresponding height and direction measurements  $h_j$  and  $v_j$ . Thus, the signal processing component fuses data from multiple range finders in a single doorway but does not fuse data from multiple doorways.

Four main signal processing algorithms are used to address the challenges. The first algorithm segments the continuous data streams into discrete doorway crossing events (Section 4.2.1). The second algorithm filters doorway events that were likely to be caused by noise (Section 4.2.2). Given



a final set doorway events, the second and third algorithms estimate the height and walking direction of the person who triggered the event (Sections 4.2.3 and 4.2.4).

### 4.2.1 Doorway Crossing Detection




Figure 4.5: The noise filtering algorithm removes baseline values from window  $\omega^i$  if they cluster with values from window  $\omega^j$ .

When a person walks under a PING range finder, one of three things will happen. First, if the sensor has already transmitted an ultrasonic pulse and is waiting for the response, the person may block the acoustic energy from reflecting back to the transceiver. In this case, the PING waits at most 18.5 msec before it times out, cancels the ranging process, and transmits another ultrasonic pulse. This is a “*timeout*” event, and it can be detected because the PING modules output a consistent max-value ranging estimate of 285 cm. This distance is typically much farther than the distance to the floor, so the resulting height estimate  $h = d_f - d_p$  is a large negative value. Second, the person sometimes blocks the line-of-sight path between the range finder and the floor, but acoustic energy still reaches the transceiver through a multi-path reflection. For example, the energy may go around the person after reflecting off the floor by bouncing off the side of the doorway and back to the transceiver. This is a *multi-path* event, and it can be detected by a small negative height estimate. Third, if the sensor has not yet transmitted an ultrasonic pulse when the person enters the doorway,




Figure 4.6: A person in the doorway can cause three different responses from the range finder: (1) Timeout events, (2) Multi-path events, and (3) Measurement events.

the next pulse will reflect off the person and return to the transceiver. An example of each of these three types of doorway crossing events is depicted in Figure 4.6.

Doorway crossings are detected by scanning all data streams in  $X$  for any timeout, multi-path or measurement event. If any such event is found, a *detection event*  $Y$  is created that contains all measurements near the event from all streams in  $X$ . More formally, any height estimates that are negative or in the range 137-198 cm (4 ft, 6 in - 6 ft, 6in) are identified. An example of such values are depicted by  $x$ 's and  $o$ 's in Figure 4.5. The timestamps of all such height estimates are clustered using the DBSCAN clustering algorithm, resulting in a set of timestamp clusters  $C$  (the timestamp clusters are not shown in Figure 4.5; the clusters shown in that figure are height clusters, as discussed in the next section.). Each timestamp cluster  $c_i \in C$  is a set of one or more event timestamps within 400 msecs of each other, denoted  $c_i = c_i^1, c_i^2, \dots$ . A time window  $\omega_i$  is defined around each cluster using the range  $\omega_i = [\min(c_i) - 200\text{msec}, \max(c_i) + 200\text{msec}]$ , and a new detection event  $e_i \in E$  is defined to be all sensor data from that range:  $e_i = X_{\omega_i}$ . A window  $\omega_i$  is illustrated by the inner bars in Figure 4.5: the event  $e_i$  includes all values within that window.




Figure 4.7: Nearby objects reflect acoustic energy and cause random noise at about 180 cm, similar to the height of one of the test subjects.

### 4.2.2 Noise Filtering

Ultrasonic range finders can be noisy, and a single outlier value will cause the algorithm above to produce a false detection. The effects of outliers are reduced by ensuring that every detection event has at least two values that are substantially different from the baseline measurements. More formally, a 200 msec region immediately before and after detection event  $e_i$  is defined, and any values in  $e_i$  that are similar to the values in that window are removed. Any outlier detection algorithm could be used here, but a clustering algorithm seems to work well: a second time window is defined that extends before and after each detection event by 200 msec:  $\hat{\omega}_i = [\min(e_i) - 200\text{msec}, \max(e_i) + 200\text{msec}]$ . A window  $\hat{\omega}_i$  is illustrated by the outer bars in Figure 4.5. All sensor values in  $\hat{\omega}_i$  are clustered, and all values from  $e_i$  that are clustered together with values in the region  $\hat{\omega}_i - \omega_i$  are removed. In the example shown in Figure 4.5, the data in window  $\hat{\omega}_i$  produces three height clusters: 1, 2, and 3. Only cluster 3 includes values in window  $\hat{\omega}_i$  but not in window  $\omega_i$ . Therefore, any values in cluster 1 are removed from  $e_i$ , leaving only the measurement events ( $x$ 's), the timeout events ( $o$ 's) and the two values in cluster 2 remaining in  $e_i$ . If detection event  $e_i$  has fewer than 2 values after removing the baseline values, then it is eliminated.

In addition to random noise, the ultrasonic range finders are also subject to periodic noise due to reflections from nearby environmental features, such as door trim, a hinge, or a nearby shelf. Figure 4.7 depicts data from one doorway sensor that frequently picked up reflections from a nearby object and produced height measurements similar to a true person's height. The key aspect of these environmental features is that they are consistently the same value and are never the negative height estimates cause by a timeout or multi-path event. Therefore, a third window is created that extends 30 seconds on either side of the detection event  $e_i$  and is used to filter these values. Any height measurement in  $e_i$  that is positive and identical to a measurement in this one minute window is removed from the event.




Figure 4.8: Incorrect height readings can be caused by multi-path reflections between ranging modules used in the same doorway.

### 4.2.3 Height Estimation

Once the detection events have been created, the height of the person that created each event is estimated. The best estimate of a person's height is typically the maximum height measurement within the event because it is taken when the person is closest to the doorway: as a person moves away from the doorway, the distance to the range finder increases and the person appears shorter. However, multi-path reflections would sometimes cause the maximum measurement to be taller than the person's true height. This is illustrated in Figure 4.8: the ultrasound signal from one range finder

reflects off the head, which triggers the other range finder to begin taking a measurement. At that very moment, a multi-path reflection (perhaps off the shoulder) reaches the second range finder and causes it to estimate a short distance (tall height).

Typically, multi-path errors only occur once in a single detection event. To eliminate these errors, the maximum height in the maximum cluster of the event is used. In the event illustrated in Figure 4.5, the maximum cluster is Cluster #3. This approach would ensure that outlier height measurements would not be used. However, this approach fails when a person walks very quickly, and the range finders only get a single reading of a person's true height. Therefore, the final height estimation algorithm is: set the height estimate  $h_i$  to be the maximum height of the maximum cluster in  $e_i$  unless there is no maximum cluster within the human height range, in which case it is set to be the maximum height measurement in  $e_i$ . The final height estimate is combined with the median timestamp from  $e_i$ , called  $t_i$  to form a new doorway event  $d_i = (t_i, h_i)$ .

#### 4.2.4 Direction Estimation

Three algorithms are used to estimate direction, each of which tries to exploit the same intuition: when a person enters a doorway, the highest heights will be measured early in the event, and when a person exits through a doorway, the highest heights will be measured later in the event. The definition of *enters* and *exits* is defined for each doorway based on the tilt angle of the sensors. The main differences between the algorithms involve how they deal with noise and outliers. Each algorithm votes on each event  $e_i$  with a +1 to declare an enter event, a -1 to declare an exit event, or 0 if it cannot decide. A direction value  $v_i$  for event  $e_i$  is defined to be the sum of the votes from all three algorithms. Because there are three algorithms,  $-3 < v_i < 3$  is the range of potential outputs.

The first algorithm uses robust least squares to fit a line to the height measurements in the event. If the slope is positive, it declares an enter event, otherwise an exit event. This algorithm only uses values above 140cm to prevent outlier values from affecting the regression. The second algorithm finds the timestamp  $t_{max}$  of the maximum height value in  $e_i$  and compares it to the median timestamp  $t_i$ : if  $t_{max}$  occurred later than  $t_i$  then the event is declared an enter event, otherwise an exit event. The third algorithm finds the timestamp  $t_{min}$  of the minimum height value in  $e_i$  and compares it to the median timestamp  $t_i$ : if  $t_{min}$  occurred earlier than  $t_i$  then it declares an enter event, otherwise an exit event.

## 4.3 Evaluation

The evaluation of signal processing algorithms explores the various components and evaluates the effectiveness of the hardware at measuring events. These two components cannot be separated and must be evaluated together. First, an experiment is defined that will be used throughout this dissertation. Second, three different components of the processing algorithms are evaluated: event accuracy, height accuracy, and direction accuracy.

### 4.3.1 Experimental Setup

The deployment consisted of 43 sensor platforms produced and placed on doorways in four different houses. Eight doorways in one of these houses are utilized for this experiment because of the difficulty in collecting baseline ground-truth data about the occupants. A sufficient long-term ground-truth system for the other houses is necessary for longer experiments to be run. Data is logged wirelessly from each doorway to a laptop located in each house and uploaded to a central server at the University of Virginia for analysis. This infrastructure is described in more detail in Chapter 6.

Three people of different heights (180 *cm*, 166 *cm*, and 160 *cm*) were utilized as test subjects. They paired up for three experiments consisting of two people each. This setup allowed testing based on people of different heights walking around simultaneously. During the experiments, each person recorded at least 500 doorway events while moving in random paths throughout the house, generating over 3,000 unique events for analysis.

Ground truth was collected using a custom-designed system that consisted of a hand-held, touch-screen interface showing an image of the house floor plan. As a person transitioned into a room, he or she would touch the destination room to indicate the room transition. The system recorded the exact timestamp, the user, and the room ID. Because this system used a touch-screen, people occasionally touched the incorrect room by accident, particularly when trying to touch the hallway because it is a small room. These events were easily found and manually cleaned from the ground truth data set because they resulted in illegal paths. The hand-held interface provided visual feedback about the room that was touched and in the case of a mis-touch, the person immediately touched the correct room before performing a new room transition. Therefore, the final ground truth data sets consisted of legal and continuous paths. A total of 20 mis-touches is estimated to have occurred out of the total 3,000 doorway events recorded.




Figure 4.9: Numbered circles indicate the location of doorway sensors and solid circles indicate the location of Motetrack beacons.

To compare with an existing solution, the Motetrack radio-based localization system [67] was simultaneously deployed in the house. Motetrack relies on a set of beacon nodes placed throughout the space that periodically transmit messages at varying transmission powers. The person being tracked must carry a battery-powered wireless receiver in order to be localized. Prior to running the experiments, 15 Motetrack beacons were deployed throughout the home, as illustrated in Figure 4.9. Each beacon was powered by a AC-powered USB adapter and were plugged directly into AC power outlets. Beacons were placed along the boundaries of rooms, near doorways and walls, although the number and locations of the beacons were constrained by the locations of power outlets. An extension cord was used to power the node in the hallway. After the experiment, the Motetrack location estimates were converted into room locations and room transitions.

### Evaluation Metric

Motetrack and the signal processing algorithms produce a sequence of events consisting of room transitions, and each transition is assigned to a test subject. To determine tracking accuracy, these sequences of events are compared to the room transitions as recorded by the ground truth system.

However, the timestamps on these two event sequences will typically not match up perfectly due to a test subject's imperfect timing between walking and touch-screen recording, touch-screen mis-touches and re-touches, and wireless transmission delays. Therefore, a one-to-one correspondence between the empirical and ground truth data is defined using a max-weight bi-partite matching algorithm. Events are only matched if they include the same person and doorway, and the weight to match each pair decreases linearly as the difference in timestamps increases, up to a maximum difference of 10 seconds. The match produced by this algorithm is illustrated in Figure 4.10, where the top row represents empirical events and the bottom row represents ground truth events. The red and blue boxes indicate whether the event is assigned to person *A* or *B*. This figure highlights four different types of matching results:

- **Type 1:** A successful match between the empirical data and ground truth.
- **Type 2:** A failed match because the tracking algorithm assigned the room transition to the wrong person.
- **Type 3:** A failed match due to a false room transition.
- **Type 4:** A failed match due to a missed room transition.




Figure 4.10: Bi-partite matching is used to address timestamp mis-matches between observed events and ground truth.

### 4.3.2 Results

The results are an evaluation of the hardware and event extraction algorithms running as a system. It is difficult to evaluate each component independently because the algorithms were designed for a particular hardware platform. *Precision* is the number of false detections divided by the number of total detections,

$$precision = \frac{\text{detected events} \cap \text{ground truth events}}{\text{detected events}} \quad (4.1)$$



and *recall* is the number of missed detections divided by the number of true doorway crossing events.

$$recall = \frac{\text{detected events} \cap \text{ground truth events}}{\text{ground truth events}} \quad (4.2)$$

### 4.3.3 False Detections and Missed Detections

| Doorway | Precision | Recall |
|---------|-----------|--------|
| 1       | 0.933     | 0.842  |
| 2       | 0.957     | 0.975  |
| 3       | 0.959     | 0.927  |
| 4       | 0.809     | 0.958  |
| 5       | 0.927     | 0.973  |
| 6       | 0.996     | 0.979  |
| 7       | 0.968     | 0.983  |

Table 4.1: The system produces more than 93% precision and 96% recall for most doorways.

Table 4.1 indicates the precision and recall for each doorway, where *precision* is the number of false detections divided by the number of total detections, and *recall* is the number of missed detections divided by the number of true doorway crossing events. Most doorways had precision levels above 93 percent and recall numbers above 96 percent. A few exceptions include doorway 4 with 81 percent precision and doorway 1 with 84 percent recall. Doorway 4 causes more false events because a shelf at roughly the same height as a person is picked up by the device. Doorway 1, Figure ??, sensors must cover a larger doorway and are mounted approximately 50 *cm* higher than the rest of the doorways. This results in a lower sample rate, thus causing more events to be missed. The signal processing achieves an average of 4 percent false negatives and 7 percent false positives across all doorways.

### 4.3.4 Height Measurement Accuracy

Figure 4.11 shows the range of heights measured for each person. The modes of these distributions roughly correspond to the differences in the subject's actual heights: 6 *cm* and 14 *cm*. However, there is substantial overlap between the full range of measurements observed for each subject. Most measurements fall within a range of 172-195 *cm*, 160-184 *cm*, and 157-178 *cm* for subjects 1, 2, and 3, respectively. This overlap indicates that many doorway events will be ambiguous and that height alone is not sufficient to identify who is walking through a doorway. This is particularly




Figure 4.11: The range of height measurements from different people have large overlap, particularly between persons 2 and 3.

true for subjects 2 and 3 that have similar heights. Height errors are not normally distributed, and measurements are more likely to be produced slightly above the mode than below. This error can be caused by many factors, including multi-path errors, as described in Section 5.2.

#### 4.3.5 Direction Measurement Accuracy

| Measure | Correct | Incorrect | Percent |
|---------|---------|-----------|---------|
| -3      | 881     | 133       | 86.9%   |
| -2      | 0       | 0         | -       |
| -1      | 190     | 133       | 58.8%   |
| 0       | -       | -         | -       |
| 1       | 230     | 84        | 73.3%   |
| 2       | 0       | 2         | 0%      |
| 3       | 714     | 111       | 86.6%   |

Table 4.2: Three algorithms produce direction measures between -3 and 3. The final direction accuracy is 81 percent across all doorways.

Table 4.2 shows the range of directions computed for each doorway event. Three different algorithms are combined to form a score between -3 and 3 with a negative value meaning a person is leaving the doorway and a positive value indicating entering. Additionally, 0 is an unknown state

where the algorithms have insufficient data and do not assign a direction. Intuitively, the more extreme a measure, the better it will represent the true direction. This is supported by the data where the -3 and 3 measures have approximately 87 percent direction correctness. These results drop the closer 0 zero the measure becomes. When the measure is 1 or -1, its correctness drops to between 58 percent and 72 percent. Very few measures occur at 2 and -2 because a single algorithm rarely fails to determine a direction when the other two compute something. The final direction accuracy is 81 percent across all doorways.

## 4.4 Conclusion

Signal processing involves different stages, which are all designed to convert raw data into events. Doorway crossing detection determines when a person has crossed through a doorway threshold. The next stage, noise filtering, reduces the effects of outlying samples and removes the various noises from the signal. Height is then estimated by utilizing the maximum height within the event even though that may be an erroneous value occasionally. Finally, direction is estimated through three algorithms.

Hardware and signal processing algorithms are able to detect an average of 96 percent of the events during the experiments. Signal processing achieves a 7 percent average false positive rate. Finally, height data from noisy measurements can distinguish between 50 and 95 percent of the data depending on the height of a person.

## Chapter 5

# Tracking

This chapter shows that the height and direction generated by our hardware and signal processing algorithms can be used to correctly assign people to rooms more than 90 percent of the time. This chapter discusses the challenges associated with utilizing a noisy sensor reading for the purposes of tracking. It then describes the tracking algorithm followed by the results which are examined with an analysis of the empirical experiments. This chapter is published as part of a *Sensys* 2012 paper

### 5.1 Design Goals

The primary goal of the system is *room-level tracking*: determining which room each person is in rather than his/her (x,y) coordinates within the home. It should handle noise, false positives and false negatives without requiring wearable tags or actions on behalf of the users.

While the signal processing algorithm operates on data from all range finders in a single doorway, the tracking algorithm operates on the detection events from all doorways. Thus, the sequences of detection events  $D$  produced by each doorway are all merged into a single sequence of events  $E$  that is processed by the tracking algorithm. Each event  $e_i \in E = (t_i, d_i, h_j, v_j)$  is a 4-tuple that indicates the timestamp, doorway ID, height measurement and walking direction associated with the doorway event. The goal of the tracking algorithm is to convert this sequence of doorway events to a sequence of corresponding room states  $S$  where each state  $s_i \in S = (r1_i, r2_i)$  is a 2-tuple that indicates the room locations of person  $p1$  and  $p2$ , respectively, that result from event  $e_i$ . For simplicity, the tracking algorithm is described in the context of a two-person tracking scenario, but a straightforward extension to three or more people is possible. The complexity of multi-target

tracking is well known to increase with the number of targets, but the system is expected to be installed in typical homes with 2-4 residents.

## 5.2 Approach




Figure 5.1: Doorway events can be ambiguous due to the possibility of sensor error (top row), but subsequent doorway events can help resolve ambiguities (bottom row).

The main challenge for the tracking algorithm is to decide which person caused each doorway event so that the person's track can be updated. In principle, height and direction measurements and/or path constraints should uniquely identify the person who walked through the doorway. However, this association is often ambiguous due to the possibility of height and direction errors, false detections, and missed detections. For example, assume person  $p_1$  is 160 cm and person  $p_2$  is 170 cm and the two people are in rooms 1 and 2, both adjacent to the same doorway as shown in Figure 5.1.a. An event detected in that doorway may indicate that  $p_1$  probably transitioned to room 2, but the possibility of a height or direction error makes this assignment ambiguous. Even if both people are not both adjacent to the doorway, as shown in Figure 5.1.b, the assignment is still ambiguous because of the possibility of false detections and missed detections: it is possible that  $p_1$  moved to room 2, that  $p_2$  transitioned to room 2 undetected and then moved to room 1, or that the detection is spurious and neither person moved. False detections and missed detections can also result in *sustained ambiguity*. For example, when a person moves back and forth between two rooms,

as shown in Figure 5.1.c, a single detection error could cause the tracking system to consistently place the person in the wrong room, with no chance of disambiguation.

A key insight of the tracking algorithm is that ambiguities can often be resolved by future observations. To illustrate this, a second detection event is shown in the second row for all 3 examples in Figure 5.1. The second doorway event shown in Figure 5.1.d disambiguates the first event in Figure 5.1.a:  $p2$  must have moved during the first event, because otherwise nobody would have been in room 1 to cause the second event would not be possible. Notice that the second event is able to disambiguate the first event even though it is ambiguous itself. Similar cases are illustrated in Figures 5.1.e and 5.1.f: the second detection event helps to disambiguate the first. To exploit this observation, tracking must be able to process future events without committing to decisions about prior events. To achieve this, it uses a *multiple hypothesis tracking* approach (MHT) in which multiple alternative tracks are considered simultaneously, each with different doorway assignments. The MHT is able to defer difficult assignments until a later time when disambiguating information becomes available. As new events are processed, tracks that are not consistent with the new information are evicted. The formulation of a MHT is described in the subsections below.

### 5.2.1 Models of Height, Direction, and Detection

In order to perform tracking, the MHT system must first model the conditional probabilities  $p(H|O)$  and  $p(V|O)$ : the probability of receiving a height measurement  $H$  or direction measurement  $V$  given the origin  $O$ , where  $O$  can be one of three things: Person A, Person B, or a false detection. The system must also model the probability of a missed detection  $p(H = \emptyset)$ : the probability that no doorway event is generated by the signal processing algorithms, even though a doorway crossing event did occur. In the current implementation, these models are learned during a training period where each individual walks under each doorway multiple times. For example, the observed height readings  $\hat{H}_{p1}$  originating from person  $p1$  are used to define the probability density function using frequency counting, as follows:

$$p(H = h|O = p1) = \frac{|\hat{H}_{p1}=h|}{|\hat{H}_{p1}|} \quad (5.1)$$

The height model for  $p2$  and both direction models are learned using the same frequency counting technique. Instead of a training period, the height and direction models could alternatively be defined to be parametric functions of people's true heights and the sensors' tilt angle. At this early stage in this project, however, empirical models are used for simplicity. During this training period, the

number of false detections and missed detections are counted and used to define the following models:

$$p(H = h|O = F) = \frac{\# \text{ false detections} / 62}{\# \text{ doorway events detected}} \quad (5.2)$$

$$p(H = \emptyset) = \frac{\# \text{ missed detections}}{\# \text{ true doorway events}} \quad (5.3)$$

where the number 62 is the total number of discrete positive height values that can be produced by the signal processing algorithm. In other words, actual height values observed for the false detections are assumed to be random, and the overall probability of false detections is uniformly distributed over the range between 137-198 cm (4 ft, 6 in - 6 ft, 6in).

Once the height, direction, and event detection models are created, the probability of missed detections is converted to a *missed transition probability*  $p(r_i, r_j)$ : the probability of going from room  $r_i$  to room  $r_j$  without being detected by any doorway sensors. To calculate the transition probability, The MHT uses a binary *adjacency matrix*  $A$ , where  $A(i, j)$  indicates whether  $r_i$  and  $r_j$  have an adjoining doorway. For example, the square, 4-room floor plan shown in Figure 5.1 has the following adjacency matrix:

$$M = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{bmatrix}$$

Using this adjacency matrix, Dijkstra's shortest path algorithm is used to find the *path length* between  $r_i$  and  $r_j$ , which indicates the number of doors through which a person must pass to transition between these rooms:  $l(r_1, r_2) = \text{dijkstra}(A, i, j)$ . Then, the missed transition probability is defined to be:

$$p(r_i, r_j) = p(H = \emptyset)^{l(r_1, r_2)} \quad (5.4)$$

Thus, the probability of staying in the same room without being detected by any doorway sensors is 1, and the probability of passing undetected through  $k$  doorways decreases exponentially with  $k$ .

### 5.2.2 Creating Tracks

In 2-person tracking, a *room state*  $s_i$  is defined to be a 2-tuple containing the room locations of persons p1 and p2 after the  $i$ th detection event:  $s_i = (r1_i, r2_i)$ . A *track* is a sequence of consecutive room states. Upon startup, the MHT algorithm creates a track for every possible combination of

initial room states: in a 2-person home with  $k$  rooms,  $k^2$  initial tracks are created:

$$\begin{aligned}
T^1 : & \quad (r1 = 1, r2 = 1) \\
T^2 : & \quad (r1 = 1, r2 = 2) \\
& \dots \\
T^k : & \quad (r1 = 1, r2 = K) \\
T^{k+1} : & \quad (r1 = 2, r2 = 1) \\
T^{k+2} : & \quad (r1 = 2, r2 = 2) \\
& \dots \\
T^{k^2} : & \quad (r1 = k, r2 = k)
\end{aligned}$$

For each new doorway event between rooms  $i$  and  $j$ , five new room states are possible: either person  $p1$  or person  $p2$  moved into either room  $i$  or  $j$ , or the doorway event is a false detection and neither person moved. To consider these 5 possibilities, the tracking algorithm duplicates every existing track 5 times and advances them in each of these five ways. For example, consider track that ends with state  $(r1 = 2, r2 = 4)$ . Upon detecting a doorway event between Rooms 5 and 6, this track would be converted into five new tracks:

$$\begin{aligned}
T^1 : & \quad \dots; (r1 = 2, r2 = 4); (r1 = 2, r2 = 4) \\
T^2 : & \quad \dots; (r1 = 2, r2 = 4); (r1 = 2, r2 = 5) \\
T^3 : & \quad \dots; (r1 = 2, r2 = 4); (r1 = 5, r2 = 4) \\
T^4 : & \quad \dots; (r1 = 2, r2 = 4); (r1 = 2, r2 = 6) \\
T^5 : & \quad \dots; (r1 = 2, r2 = 4); (r1 = 6, r2 = 4)
\end{aligned}$$

The number of possible tracks increases exponentially with the number of doorway events. After processing  $m$  doorway events, a total of  $5^m$  tracks will have been created. Furthermore, each track will have  $m + 1$  room states:  $T^i = (s_0, s_1, s_2, s_3, \dots, s_m)$ , where  $s_0$  is an initialization state and state  $s_i$  was created in response to doorway event  $e_i$ .



### 5.2.3 Weighting Tracks

Every hypothesized track  $T_i$  is given a weight  $\omega_i$  that is proportional to the probability of the track given the observed doorway events  $D$ :  $\omega_i \propto P(T_i|D)$ . The weight of all initialization tracks is set to 1, and when room state  $s_m$  is added to a track in response to doorway event  $d_m$ , the track's weight can be updated incrementally by multiplying by the probability of transitioning from  $s_{m-1}$  to  $s_m$  given  $d_m$ , as follows:

$$\omega_i^m = \omega_i^{m-1} * p(o_m|h_m) * p(r_{m-i}, r_p) \quad (5.5)$$

where

$$p(o_m|h_m) = \frac{p(H = h_m|O = o_m)}{\sum_{o=(A,B,F)} p(H = h_m|O = o)} \quad (5.6)$$

and where  $h_m$  is the height measurement of doorway event  $d_m$ , and  $o_m$  is the origin of doorway event  $d_m$ :

$$o_m = \begin{cases} A & \text{if } rA_{m-1} \neq rA_m \\ B & \text{if } rB_{m-1} \neq nrB_m \\ F & \text{otherwise} \end{cases} \quad (5.7)$$

and  $r_{m-1}$  is the last known room location of the event origin

$$r_{m-1} = \begin{cases} rA_{m-1} & \text{if } o_m = A \\ rB_{m-1} & \text{if } o_m = B \\ \emptyset & \text{otherwise} \end{cases} \quad (5.8)$$

and  $r_p$  is the previous room location of the event origin, which is opposite the doorway of the origin's new room location

$$r_p = \begin{cases} \text{opposite}(rA_m, d_m) & \text{if } o_m = A \\ \text{opposite}(rB_m, d_m) & \text{if } o_m = B \\ \emptyset & \text{otherwise} \end{cases} \quad (5.9)$$

Because every possible explanation for doorway events  $D$  is considered, the weight of each track  $T_i$  is proportional to its probability  $\omega_i \propto p(T_i|D)$ , and could be converted to that probability if

normalized by the sum of all track weights:

$$p(T_i|D) = \frac{\omega_i^m}{\sum_j \omega_j^m} \quad (5.10)$$

Because the denominator is the same for all tracks, the weight alone is sufficient to produce a global ordering of all tracks.

#### 5.2.4 Merging and Evicting Hypotheses

The number of tracks under consideration grows exponentially with the number of doorway events detected, and will quickly exhaust memory and computational resources. To address this problem, a “ $n$ -best” eviction policy is used to eliminate the least likely tracks. After each doorway event is detected, the MHT temporarily increases the number of tracks to  $5n$ . These tracks are weighted, sorted, and the  $4n$  tracks with the lowest weights are evicted, as illustrated in Figure 5.2. Thus, the number of tracks under consideration is limited to a constant number  $n$  after each time step.

The main challenge for this eviction policy is that ambiguous doorway events cause existing states to be replicated many times, but do not help differentiate the replicas and so many of them have approximately the same weight. As more ambiguous events are encountered, these replicas quickly multiply and dominate the track buffer, causing viable tracks to be evicted. Because of the 5x multiplier for each detection event, the “ $n$ -best” eviction policy allows the MHT to consider approximately  $\lfloor \log_5(n) \rfloor$  ambiguous doorway events before it makes arbitrary assignments due to space limitations. Even at values of  $n = 1000$ , MHT can only fully consider 4 ambiguous doorway events at a time.

In order to consider more ambiguous events simultaneously, a track merging algorithm was developed that finds and eliminates unnecessary replicas. The basic intuition is that future information will no longer help resolve ambiguous doorway assignments that occur before a fully known state. For example, if after event  $m$  all tracks agree on the state for step  $m - 1$ , no new information can help resolve ambiguity about the states prior to  $m - 1$ . Therefore, MHT can resolve multiple hypotheses about any event prior to the known state. After the eviction step, MHT follows a 4-step merging procedure: (1) it scans for any event about which all tracks share consensus (ii) it *commits* to the values of the track with the highest weight for all states  $m - 1$  and earlier, i.e. it outputs these state values to the user (iii) it truncates all tracks at step  $m - 1$ , as illustrated in Figure 5.2, and (iv) it searches all remaining tracks for any exact duplicates: tracks that have exactly the same set of state

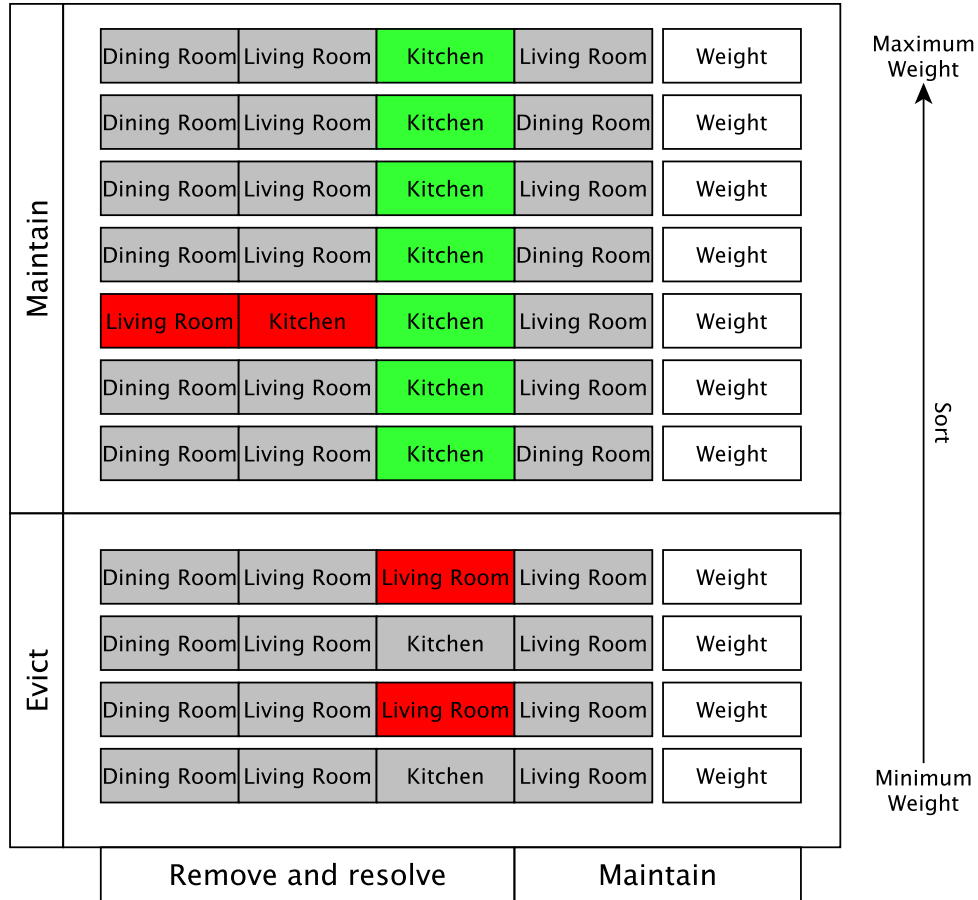


Figure 5.2: The MHT is composed of many tracks (rows). These are sorted according to their weights and a portion are evicted. The third column is resolved when all tracks agree on the event assignment. Tracks are resolved by removing all preceding state and shortening each individual event sequence.

assignments, and were only differentiated by the states that were just committed. For every pair of duplicate tracks  $i$  and  $j$ , duplicate  $j$  is evicted and duplicate  $i$ 's weight is updated to be:

$$w^i = w^i + w^j - w^i * w^j \quad (5.11)$$

This 4-step merging algorithm evicts ambiguous events that will never be resolved and simultaneously reduces the length and number of tracks under consideration. By doing so, it makes space in the track buffer for new ambiguous events to be considered.

## 5.3 Simulation-based Evaluation

A tracking simulator was built that simulates people moving between rooms within a house based on the following parameters:

1. Simulation length.
2. A sensor noise model: including mean and standard deviation for a Gaussian noise model and the probability of a timeout event where there is no height information.
3. The heights of people.
4. The probability of false negatives or false positives.
5. Timing variance which is the difference between event and ground truth times.
6. The floor plan.

This simulator generates input and testing files that are identical to the empirical tracking implementation so that the same tracking algorithm code can be applied to the simulator output as the hardware output. Each simulation was run an identical number of times which results in approximately the same number of events for each person across all trials. The exact number of events vary based on the number of false positives, false negatives, and general randomness of the simulation.

### 5.3.1 Experimental Setup

Tracking accuracy is the primary evaluation metric and is computed for each person across all events. *Accuracy* is the number of room transitions correctly identified by examining room transitions for each person and validating with the ground truth data. Examining each transition independently allows the ability to prevent errors in one assignment from affecting the state of the other person. Room transitions provide a better metric for tracking comparison because they remove the time aspect and only check when people are transitioning. For example, a person that entered and stayed in a single room could have a great tracking accuracy even if a large number of short events were all classified incorrectly. However, an evaluation based on room transitions tests the ability of the algorithm to track someone throughout any period of time. These two evaluation methods - transitions and rooms - can easily be converted to each other but due to the desire to evaluate directly tracking, room transitions are used in this dissertation.

Simulation is used for an extensive sensitivity analysis that could not be performed in a real deployment. First, an evaluation of the false positive and false negative rates show how they affect accuracy. Next, sensor noise and the number of sensor timeouts are adjusted to test the reactions of the MHT to imprecise data. Finally, an evaluation shows the effects of varying the height difference between two people and the effect of having a correct direction on tracking accuracy. Second, a real-world evaluation is used that simultaneously evaluates all components in a house. Additionally, variations of the house topology are tested.

### 5.3.2 False Positives

The first experiment tests how the false positive rate affects tracking accuracy. The simulation varies the false positive rate between 0 and 70 percent where a 100 percent false positive rate indicates an equal numbers of errors and events. Figure 5.3 illustrates how false positives affect tracking accuracy (score), the uppermost line indicates recall, and the lowest indicates precision, and the middle line shows the score which is computed by taking the harmonic mean (F-Measure).

$$\text{F-Measure} = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \quad (5.12)$$

Two results are apparent in this figure. First, the tracking algorithm is operating properly with 100 percent tracking accuracy when there are no false positives. Second, precision and recall drop as the false positive rate increases because once the false positive rate saturates the ability for the MHT to resolve state, it starts getting tracking data incorrect. Tracking accuracy is expected to exceed 95 percent based on the empirical false positive rates of seven percent on average from Section 4.3.3.

### 5.3.3 False Negatives

False negatives are instances where a doorway event is not detected. Increasing the false negative rate has a much larger impact on tracking accuracy than the false positive rate. Figure 5.4 shows the same three lines as Figure 5.3 except precision is on top and recall on the bottom. False negatives impact tracking accuracy in a nearly linear manner. Empirical experiments in Section 4.3.3 indicate false negative rates between 1.5 and 15 percent for the hardware and software configuration. This results in an expected tracking accuracy above 93 percent. A nearly linear drop in tracking accuracy is expected because the evaluation metric compares room transitions and the tracking algorithm is unable to test a state that is missing. Precision also drops as a result of varying false negatives




Figure 5.3: This figure shows the effect of increasing the false positive rate while maintaining zero false negatives. There is a nonlinear relationship between the false positive rate and tracking accuracy.

which is due to the tracking algorithm attempting to resolve paths into consistent states by marking some as false positives.

### 5.3.4 Sensor Noise

Sensor noise occurs for any source of measurement problems not removed by signal processing. Figure 5.5 models this noise as a Gaussian distribution centered at the height of each person, 166 *cm* and 180 *cm* respectively. This analysis utilizes a false positive rate of 9 percent and a false negative rate of 5 percent, and by adjusting the variance between 1 *cm* and 26 *cm* the amount of noise is tested. Figure 5.6 illustrates this trade-off and the linear relationship between tracking accuracy and the amount of noise. Precision and recall both lose the same amount of accuracy as noise increases. This result is determined by the number of consecutive and non-ambiguous paths each person takes. Since the test house has a limited number of rooms to utilize, the tracking accuracy drops as a function of sensor noise.




Figure 5.4: This figure shows the effect of increasing the false negative rate while maintaining zero false positives. There is a nearly linear relationship between false negatives and tracking accuracy.

The second part of sensor noise is tested by varying the probability of timeout events from 0 to 100 percent. The simulator has a noise model matching measured data and by fixing the false positives and false negatives rates the affects to tracking accuracy are shown. Figure 5.7 shows the effect of increasing the number of events with timeout values on tracking accuracy. Tracking accuracy drops from 94 to 88 percent during the first 50 percent of inserting timeout events; however, raising the timeout rate further causes this accuracy to drop quickly to 50 percent, which is due to the tracking algorithm guessing randomly at person assignments since there are not enough identifying measurements for it to function properly. Empirical experiments indicate that approximately 10 percent of the events will contain timeout measurement which produces a tracking accuracy of more than 90 percent.

### 5.3.5 Occupant Height

The difference in heights between people affects tracking accuracy. It is one thing to track with a height difference of 30 *cm* and another when it is 1 *cm*. The simulator was used to evaluate this height difference. The empirical experiments used test subjects that were 160, 166, and 180 *cm* tall.




Figure 5.5: Gaussian Noise examples that illustrate an overlap in measurements for two different people.

The simulator varied accordingly by setting one person at a static height of 160 *cm* and covering the range between 161 and 189 *cm* for the other person. Figure 5.8 shows a linear relationship between the height difference and tracking accuracy. This figure confirms the intuition that larger differences in height result in better tracking results. Precision and recall score above 90 percent with a large height difference; however, when the difference is a small (1 *cm*), the accuracy drops to 70 percent.

### 5.3.6 Direction

The effect of direction error is evaluated by randomly adjusting the event direction to an incorrect state based on a predefined probability. Figure 5.9 shows the effect of changing direction correctness from 0 to 100 percent. Precision varies little across the entire range, indicating that tracking is handling false positive events at similar rates; however, recall tells a different story. At either extreme, 0 and 100 percent, it performs well which indicates that path history overrides direction. Recall is lowest in the middle (50 percent) because randomly choosing a direction is detrimental to an algorithm that relies on path information. At the 100 percent correct state, accuracy is better by






Figure 5.6: Sensor noise impacts the tracking algorithm by causing imprecise readings. This figure shows how the increasing noise degrades the tracking performance from about 93 percent to below 84 percent.

6 percent than having all incorrect states which shows that direction has a small impact on tracking performance.

## 5.4 Experimental Evaluation

In order to evaluate the trade-off between computational resources and tracking accuracy, the tracking algorithm is tested while varying the number of tracks that can be considered simultaneously. The results are illustrated in Figure 5.10, and demonstrate that a MHT can store, update, and compare use few as 20 alternative tracks. While processing the 3000 doorway events in the experiments, tracks stored in memory were 24 states long on average, with a maximum length of 55 states. This indicates that the MHT can be executed on memory-constrained devices such as a cell phone without sacrificing tracking accuracy. Note that even 1000 tracks were not sufficient before we added the track merging algorithm described in Section 5.2.4.

The algorithm can also be executed in real time: using 4-core Intel i7 processor running at 2.67 GHz, all 3000 doorway events could be processed in a total of 280 seconds with 180 seconds dedicated




Figure 5.7: An important aspect to analyze is how the tracking algorithm handles the timeout noise. This figure shows how tracking accuracy changes as the amount of timeouts increases, dropping to 50 percent at the lowest.

to signal processing and 100 seconds for a Python implementation of the tracking algorithm. The MHT is an on-line tracking algorithm, meaning that it does not need to process all data in batch. However, it does require look-ahead for noise filtering and data association. The signal processing algorithm uses a 500 millisecond look-ahead window, which adds a small delay to doorway event detection. In the experiments, the MHT tracking algorithm used a look-ahead of 14 doorway events on average before committing to a doorway assignment, with a maximum look-ahead of 49 doorway events. The length of this delay depends on how frequently people in a house move between doorways, and may even be several hours. However, just like other data association algorithms, MHT can make an immediate estimate of each person's location based on the distribution of states in the tracks under consideration, and can use the variability in those states to provide a confidence value. In this usage mode, the user may see changes in the estimated location if future events cause the system to make a different decision.




Figure 5.8: The system is designed to operate with people of different heights to understand the tracking algorithm’s sensitivity to variation when faced with sensor noise. This figure shows how heights affect tracking accuracy in the presence of noise in a roughly linear manner.

### 5.4.1 Topology Changes

Once data was collected for all doors during the initial experiment, rooms corresponding to leaf nodes are removed from the topology one by one until only two rooms are left (Figure 5.11). Initially, there are five leaf nodes in the test topology. All the initial leaf nodes are removed before moving on to interior nodes. By removing the appropriate doorways before execution, the tracking algorithm simulates what would have happened if the rooms did not exist. Figures 5.12(a), 5.12(b), and 5.12(c) show a similar pattern. At first, removing rooms only slightly decreases the event recall rate. However, the tracking algorithm starts having trouble keeping track of the individual person states where fewer than four rooms being considered. All experiments drop significantly due to the lack of sequential path information which forces the algorithm to rely only on measured values. Ground truth data also drops at this cutoff point for the same reason.

The test house was chosen for two reasons: first, it contains the widest connected room topology which allows tracking to take advantage of the longest possible paths, and second, the large topology size allows it to be virtually reduced by removing rooms from the data traces to test various topologies. Removed rooms are dropped from the event trace to make it appear like a person never




Figure 5.9: Direction accuracy is tested by altering the percentage of correct readings. Tracking performs well when either all or none of the directions are correct. When a mixture occurs, accuracy drops.

entered them. These experiments show tracking accuracy drops based on the number of rooms in a house and significantly when the number of rooms reaches a critical limit which is determined by persistent ambiguity and cascading failure rates. The fewer rooms that exist, the more these errors are likely to occur.

#### 5.4.2 MHT accuracy

In order to test the effect of height differences on tracking accuracy, we performed 3 controlled experiments with different pairs of test subjects  $(p1, p2)$ ,  $(p1, p3)$ ,  $(p2, p3)$ , where each subject had a different height:  $p1 = 180\text{ cm}$ ,  $p2 = 166\text{ cm}$ , and  $p3 = 160\text{ cm}$ . In each experiment, the subjects walked randomly through the house shown in Figure 4.9 that contains 7 interior doorways. Exterior doorways were not used in this experiment. The rooms in this home were separated by at most 4 doorways, and most doorways in this home were  $90\text{ cm}$  (3 ft) wide x  $213\text{ cm}$  (7 ft) high and were covered with 2 ultrasonic range finders. Doorways (1) and (2) were  $256 \times 274\text{ cm}$  and  $121 \times 213\text{ cm}$ , respectively, and were therefore covered with 4 and 3 ultrasonic range finders, respectively. During




Figure 5.10: A 4-step path merging algorithm allows MHT to perform well with only 20-30 simultaneous tracks.

the experiments, each subject recorded at least 500 doorway events with the experiments produced a total of over 3,000 unique doorway events for analysis.

The tracking results for both the MHT and MoteTrack are shown in Figure 5.13, where each set of 3 bars represents the results for the three experimental runs. The 1st set of bars shows that Motetrack averages 37 percent room-level tracking accuracy. In comparison, the 2nd set of bars shows that the MHT achieves 90 percent room-level tracking accuracy on average.

Motetrack's performance can be explained by several factors. First, Motetrack performance increases with the number of beacon nodes, and its poor performance in these experiments is partially due to the fact that 15 beacon nodes were used (roughly 2 per room). Second, the Motetrack receiver was carried in pant pockets throughout the experiments. Therefore, rotation of the body would cause significant changes to the radio signal strength measurements, appearing as physical movement even if the person was only turning in place. Most importantly, however, Motetrack is not designed for room-level tracking: it locates a person in space and small deviation in (x,y) coordinates can lead to large errors in room location and room transition, by causing a location estimate to be on the wrong




Figure 5.11: The floor plan of the test house. The numbers indicate which rooms are part of each topology. If there are 8 rooms in the topology, then rooms numbered 1 to 8 are included.

side of a wall or doorway. The results of this experiment do not demonstrate that MHT is a better tracking system than Motetrack, but only that it achieves higher room-level tracking accuracy, when used with approximately 2 beacon nodes per room.

The extent to which MHT accuracy is affected by false detections, missed detections, height errors, and direction errors is illustrated in Figure 5.13. The 3rd set of bars shows that tracking accuracy improves to 90.4 percent on average when false detections (Type 4 matches) are removed from the empirical data set. The 4th set of bars show that accuracy improves to 95 percent on average when ground truth height and direction values are used. The 5th set of bars shows that accuracy improves to 99.9 percent when missed detections (Type 3 matches) are replaced with ground truth events. This is equivalent to running the tracking algorithm on the ground truth data set, and illustrates that the tracking algorithm itself is not introducing errors. From this analysis, false detections, height/direction errors, and missed detections account for 0.4, 4.6, and 4.9 percent of tracking accuracy, respectively.


## 5.5 Conclusion

Simulation is a practical method for extensively testing the numerous parameters of the tracking algorithm and quickly generating many different scenarios without needing to enter a house and perform extensive data collection. It also allows some parameters to be varied that are impossible or require complete hardware redesigns, such as sensor noise, accuracy, or the signal processing accuracy. Real-world data collection is also a poor choice when attempting to evaluate the effects of tracking performance on people of various heights. It is difficult and time consuming to recruit 30 different people all varied by 1 *cm*. Simulation has the added benefit of generating both the input files for tracking and the ground truth files used to evaluate the algorithm's performance, thus eliminating any source of error associated with ground truth generation.


The MHT addresses two of the three challenges associated with simultaneously identifying and locating two people in a house. First, it addresses cascading failures by including a weighting parameter that accounts for the physical topology of the house and utilizing successive measurements to have more confidence in the results. Teleportation is addressed by these same sequential measurements and utilizing false positive and false negative rates in the weighting function. Finally, an attempt is made to address persistent ambiguity by utilizing the sequential measurements and maintaining various hypotheses; however, unlike the continuous domain with a lot of opportunities for resampling, the houses have less than 10 rooms, which results in paths that continually overlay. This creates cycles that any tracking algorithm cannot address. The correct way to fix persistent ambiguity is to address it at the hardware layer where better data can be provided about each event.

Simulation allows the effects of false positives and false negatives on the tracking algorithm to be tested by varying sensor noise and testing people of different heights. By setting all simulation parameters to those that reflect real-world experiments evaluated, the tracking algorithm is evaluated with a realistic configuration and the results approximate the empirical values.


Empirical results from three experiments result in a 90 percent tracking average. MHT perform comparably to the hypothetical alternative which removed all false positives from the data set indicating it is robust to the amount of noise generate by the custom sensors. It was never expected to achieve the accuracy when false negative data points were handled. When compared to the Motetrack system, MHT outperforms it by at least a factor of two. This is especially important considering MHT utilizes no wearable technology that must be maintained.



(a)



(b)



(c)

Figure 5.12: House topology was altered to reduce the number of nodes in the house and measure the recall of tracking for all experiments.






Figure 5.13: Motetrack achieves 37% tracking accuracy and the MHT achieves an average of 90% accuracy. Analysis indicates how much results improve when false detections, height/direction error, and missed detections are removed.

## Chapter 6

# Deployment

This chapter discusses the challenges of designing, deploying, and managing large-scale residential sensing systems such as the ones presented in Chapters 3, 4, and 5 and is published in *Sensys* 2011. 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 custom designs. Each of these deployments was performed for a different application or experiment, and all deployments review at once to distill out a single set of guidelines and design principles supported by data, pictures, and anecdotes from the deployment experience. By focusing on common myths and misconceptions, future deployments may avoid potential pitfalls. Most of the guidelines identified are common sense, but are obvious only in hindsight; the contribution of this chapter is the collection of hazards and challenges that make these guidelines appear obvious.

The key finding of the 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. The scale of residential sensing deployments was pushed 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. 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 encountered indoors are identical to those already observed in outdoor deployments [68]. Other challenges however, were entirely unexpected or were analogous but not identical to those out-

doors. The deployment experience helped to elucidate these similarities and draw analogies where possible. Some surprising findings were identified as the deployments scaled and 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 a greater number of different COTS platforms were used, because each new platform introduced a new set of possible failures.

## 6.1 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, the deployments in this dissertation 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 therefore they are expected 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). The deployments in this dissertation 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 [69] 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 [70] 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 [71] 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 [72]. "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 [73]. Cook et al. deployed 18 motion sensors and 2 temperature sensors for four months [74]. Kasteren, et al deployed 14 state-change sensors for four weeks [75]. <http://boxlab.wikispaces.com/List+of+Home+Datasets> 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. Section 6.4, describes several reasons why these systems are difficult to scale to hundreds of nodes.

**Lessons learned:** Many of the lessons discussed in this chapter 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 chapter does. For example, Edwards and Grinter discussed seven challenges that must be overcome in order for smart homes to be viable [76], including possible social ramifications, the effect on home life, the challenges 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” [77] 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, this chapter 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” [68]. 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 chapter, the indoor deployment are reflected on to identify similarities and draw analogies to outdoor deployments when possible in order to

|   | #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 6.1: Over the course of several years, 100's of sensors in over 20 homes were deployed, ranging from small, short-term to large-scale and long-term deployments. This table summarizes the main deployments. One home was used for deployments A, B and H, and another home was used for both deployments F and G.

strengthen and generalize the lessons learned from both, even if they are not obviously comparable at first.

This chapter is the first compilation of potential pitfalls, common misconceptions, and practical advice about indoor sensing in residential environments.

## 6.2 Overview of the Residential Deployments

Over the past several years, over 1200 sensors in over 20 homes have been deployed 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. The 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 [22,58,78,79], 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 dissertation. In this section, the sensors and the system architecture that was deployed in order to provide more focused context and perspective for the analysis and design guidelines that follow is described.

**Sensors and Controllers:** The sensors and actuators in the deployment changed with each home, each experiment, and each generation of hardware. A summary of the main deployments is shown in Table 6.1, and images of some of the sensors are shown in Figure 6.1. 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 6.1(d)). Motion sensors manufactured by X10 were used 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.

Contact reed switches (latch sensors) were used to detect the use of objects, including appliances, cabinets, doors/windows, lights, and water fixtures (Figure 6.1(h)). X10 sensors were originally used due to cost. In later experiments when the contact switches served as ground truth sensors, however, Aeon Labs Z-wave devices that use a reliable communication protocol provided better data. Reed




Figure 6.1: A wide array of sensors in homes were deployed, a subset of which are shown in this figure. Commercial products were used whenever possible and designed and integrated custom solutions when necessary. Latch sensors designed for doors and windows were used to detect use of other objects in the home, including appliances, cabinets, light switches, and water fixtures (h). Multiple generations of hardware designs for several of the sensing sub-systems were used, 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).



switches were too obtrusive to use on light switches for long-term deployments, and so multiple iterations of a custom design based on the synapse-wireless SNAP mote, and off-the-shelf GE Z-Wave home automation switches (Figure 6.1(a)) were tested.

Ultrasonic range finders were deployed 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 6.1(c)). To evaluate MHT, wearable RF beacons and the motetrack tracking software [29] were used.

The TED 5000 was used for whole-house power metering, the Powerhouse Dynamic eMonitor for circuit-level monitoring (Figure 6.1(f)), and an Aeon labs Z-wave meter for plug load monitoring (Figure 6.1(b)). Shenitech's ST301 transit-time ultrasonic flow meter was used for water metering and hot water usage monitoring (Figure 6.1(g)). Different platforms were used 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.

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

**System Architecture:** The 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 6.2 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 the sensing sub-systems were robust to both short-term power outages and broadband disconnections. A 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 at the University of Virginia.

Many of the sensing sub-systems did rely on AC power at some point, and therefore lost data even during short-term power outages. For example, the 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. The 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.

Two 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.




Figure 6.2: The system architecture includes a heterogeneous array of sensing sub-systems, each of which is subject to a different set of failure modes.

### 6.3 Reliability and Failure Analysis

In this section, the main failure modes of the system during four recent deployments are examined.

**Failure Detection and Classification:** Down time intervals for each sensor are analyzed by defining the longest acceptable time interval  $\tau$  between two consecutive data points; any interval longer than  $\tau$  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 6.2 summarizes the parameter  $\tau$  used for each sensor type.

Down time intervals for each sensor are used to identify the root cause of failure. By exploiting the fact that each failure mode causes data loss in a distinct subset of sensors, the root cause of each failure can be identified based on the set of simultaneous sensor failures, as follows:

1. **Wireless link loss:** down time of a single wireless sensor for less than  $4\tau$ <sup>1</sup>
2. **Battery dead:** down time of a single battery-powered sensor for longer than  $4\tau$
3. **Plug disconnected:** down time of a single plug-powered sensor for longer than  $4\tau$
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

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

Table 6.2: Periods of down time for each sensor type are identified by defining the longest acceptable time period  $\tau$  between two consecutive data points.

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 6.1) and the power, communication, and gateway or Internet resources used by each sensor (Section 6.2).

<sup>1</sup>Wireless link loss is not assessed for event-driven X10 and Z-wave sensors because they do not transmit periodically.

Identification the root cause of failures is similar to that of Sympathy [80]. The key difference is that the 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, custom scripts and a tool called Nagios are used 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 presented here uses hindsight to improve failure analysis.

**Results:** The failure detection and classification algorithm described above are used on the four deployments named *G*, *H*, *I*, and *J* in Table 6.1 for the seven-month period from January 1, 2011 to August 1, 2011. The effect of a system failure is measured in terms of sensor down time, using a metric called *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.

Table 6.3 shows the sensor down time due to each type of system failure in each house, and Figure 6.3 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 the 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.

| 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 6.3: The total sensor down time for four of the deployments listed in Table 6.1, broken down by root cause and measured in *sensor-days*: #days \* #sensors.

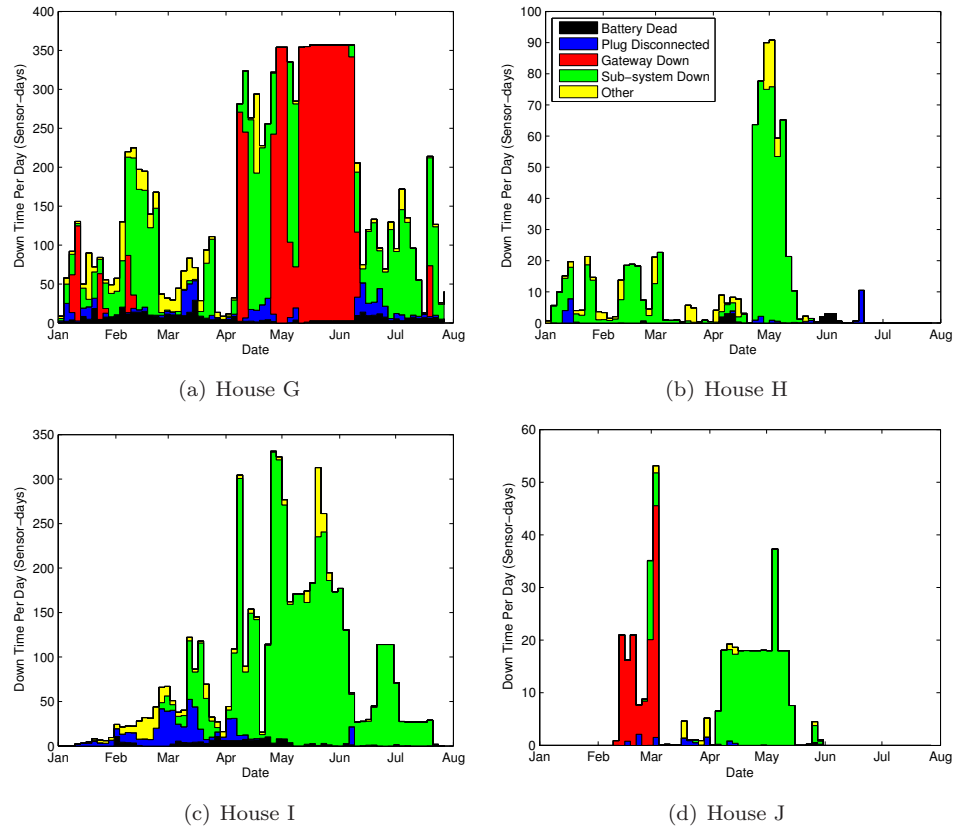


Figure 6.3: Total sensor down time for four deployments, broken down by root cause over time. Individual sensing sub-system failure is the dominant cause of down time, followed by failure of the gateway machine.

An 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 than the other two houses.

Wireless link loss is often a main challenge for reliable data collection, but in the 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 a new USB address and not be found by the software stack. Sub-system failures were so common because the system used nearly a dozen different COTS sub-systems, each of which failed infrequently but in aggregate caused substantial down time. Sub-system failures became worse in all homes after April 8th 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.

## 6.4 A Hitchhiker’s Guide

A set of guidelines was produced based on the experience and analysis of the deployments to avoid many of the observed pitfalls and failures. Future 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.

### 6.4.1 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. Three different ways of powering sensors in homes were explored and concluded that batteries still have an important place in large-scale residential sensing systems.

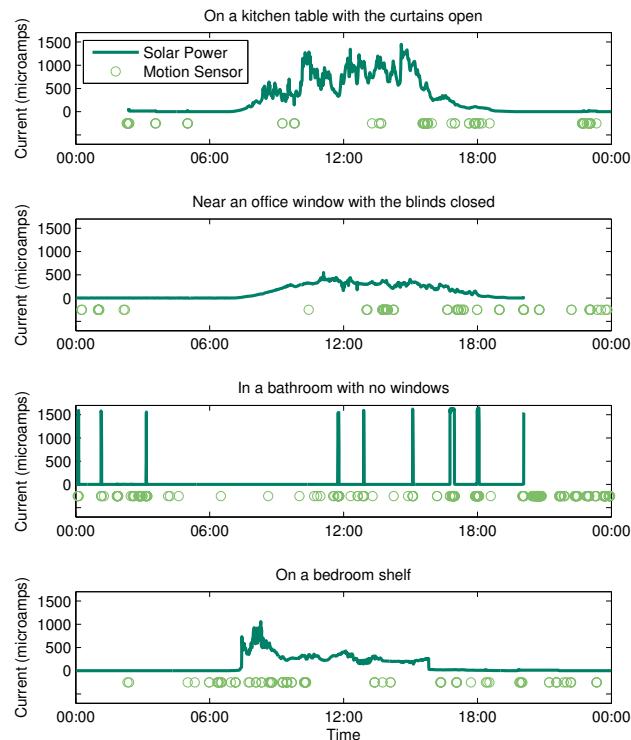


Figure 6.4: Solar energy can be harvested indoors for only 8-9 hours per day, despite much longer days and indoor light usage, and the energy harvesting time period does not necessarily correspond with times of human activity. The average power that can be harvested changes dramatically in different rooms and even locations within a room.

**Wall Sockets:** 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, users will unplug sensors for their laptops, hairdryers, or vacuum cleaners. This became salient in one deployment when all but 6 out of 40 power receptacles in a house were used for sensors, many of which were filled beyond normal capacity using expansion adapters and power strips (Figure 6.5(a)).

Furthermore, contrary to the popular belief that wired power is more reliable than battery power, in the 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 over 250 linear feet of wire was necessary 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 6.5(b)) or near frequently moving objects (Figure 6.5(c)). 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, approximately two dozen service calls were made 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 to use batteries for the second generation light switch sensor (Figure 6.1(a)): as an experimental prototype, the software is expected 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 analyzed, two of the 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 the 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.

| 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 6.4: The devices used in the deployments have a wide range of average and peak power characteristics.

**Indoor Solar:** Indoor solar power is one alternative to wall power and battery power. Figure 6.4.1 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. 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 6.4). For comparison, this same solar panel produced at most 102mW outdoors.

### 6.4.2 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, thirteen 802.15.4 devices were placed 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. One deployment was in a house with concrete slab flooring separating the three levels and copper siding that produced a Faraday cage, isolating wireless sensors outside. The 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. The deployments 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 narrowband 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 this was observed in the deployments with up to two weeks of almost continuous data loss on the power line communication system used by the TED power meters.

### 6.4.3 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. A number of unlikely causes of sensor failure in homes have been identified 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 6.5(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 the first-generation light switch sensors were easily snagged (Figure 6.1(a)), and the faucet sensors were subject to both wet surfaces and user interference (Figure 6.1(h)). The mounting techniques should not rely on users learning to accommodate sensors, because a large fraction of dislodged sensors in the 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 6.5(c)). Sensors installed on furniture such as bookshelves may be moved or hidden and produce changed, invalid, or unreliable data.



(a) Overloaded Sockets



(b) Snag Hazard



(c) Snag Hazard

Figure 6.5: 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).

**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 over 500 sensors were 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, a set of automated scripts was deployed 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. The 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, all critical alerts were projected onto a wall in the lab, together with the duration of the alert (Figure 6.4.4). 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 6.3.

frequency. a threshold. accessible. sensors and actuating devices are running. output

#### 6.4.4 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 [81], it is 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. 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 needed to be checked. 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 6.1(d)). Cover LEDs before deployment; unlike programmable LEDs on experimental platforms, LEDs on COTS devices (Figures 6.5(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.

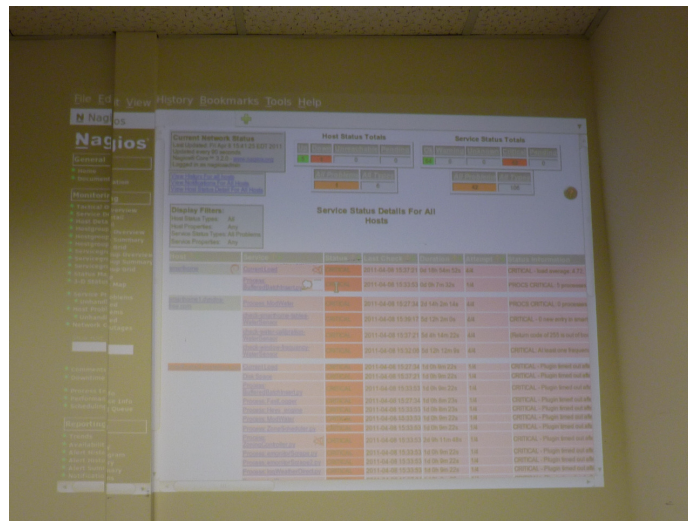


Figure 6.6: The system constantly verifies system operation and critical alerts are project onto a wall at the University of Virginia.

**Test Three Times:** Several outdoor deployment studies have emphasized the need for both lab testing and deployment time validation [68]. During the deployments, it was critical to test

immediately *after* deployment, because sensors that are designed to monitor people are greatly affected by the experimenters themselves. Thus, testing not just be twice, but three times: before leaving the lab, at deployment time, and immediately after leaving the site.

In one of the deployments, over 60 motion sensors were installed 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, the second generation doorway sensors were installed 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.

#### 6.4.5 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 [70], a human observer in the home [73], or controlled experiments with pre-determined activities [82]. 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, 5 such studies were performed, each requiring a different level of user participation, and found that data could be collected either with high accuracy or for long time periods, but not both simultaneously. In real-time tracking studies that required constant participation, high quality data could only be achieved 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. This trend is not likely 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? The 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 consistent data could be collected 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. 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 the studies required only 2 sensors at the electrical and water mains of the house, but we needed to install over 100 sensors to validate the measurements by monitoring all light switches, plug loads, faucets and water fixtures, and major appliances. These object sensors can also be used 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.

### 6.4.6 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, the 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 6.1(a)). Due to push back from users, the later generation systems were designed 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 6.1(c)). The COTS light switch sensor hides the electronics behind the switch plate (Figure 6.1(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. Generous quantities of double-sided tape were used in the 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. Later systems used painter’s tape (Figure 6.1(c)), but even that peels paint when left too long. The 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, the 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 electronics devices such as televisions and alarm clocks, but become a first-class problem when scaling to hundreds of sensor nodes. LEDs on the 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 the 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 6.1(c)) or deep inside air ducts (Figure 6.1(e)) because indirect reflections cause these features to glow at night. The short-duration, event-triggered LEDs on the light and faucet fixture sensors (Figure 6.1(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 the first generation height sensors (Figure 6.1(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 the 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 uninterruptible 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 the 120VAC-to-5VDC converters (Figure 6.5(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.




Figure 6.7: As an increasing number of X10 motion sensors are added to the network, packet deliver ratio (PDR) drops from 100% to near 20%.

#### 6.4.7 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 sub-systems 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 COTS devices were used whenever possible (Figure 6.1). 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 maintenance effort increases with the number of different sub-systems more than it does with the total number of sensors.

For example, 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, the 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 [81]. 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 6.1(a)) do not have reliable transmission, 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 6.1(d) and 6.1(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. Performance was quantified in a controlled experiment with an increasing number of sensors, using five trials per configuration (Figure 6.7) and showed that PDR drops to near 20 percent 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.

## 6.5 Conclusions

After deploying over 1200 sensors in over 20 homes, numerous facts and insights about the challenges, hazards, and pitfalls of residential sensing were identified. These experiences are key for success in future deployments, and the full set of all deployments was used to distill a single set of guidelines,

each supported by data, images, and anecdotes from the experience, with the hope of preventing others from repeating these 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. These similarities are acknowledged and embraced, and the challenges of residential sensing have been articulated in those same terms. By doing so, the differences are easier to articulate: the nature of limited access, the types of hazards in the environment, and the reasons why large scale and long-duration deployments become challenging.

## Chapter 7

# Conclusion

Whether he is inside or out, man desires the ability to define where he is and use that knowledge to improve life. A compass or GPS provides that information when a person is outside, but no modern technology exists to provide it when a person's location moves indoors without interfering in day-to-day life. The system presented in this dissertation offers progress towards such a solution.

Many challenges appear when creating a system to track and identify people without their active participation. A biometric must be chosen and evaluated for its effectiveness in reliable tracking. This biometric must not require either sensors to be placed on participants' bodies or cameras to be placed within their homes. Algorithms must be created and tested for their ability to process vast amounts of data. Simulation experiments must be run to see whether placing this system in real-world scenarios is worth the time and effort. Above all, the system must be made palatable to participants; otherwise they will not allow it into their homes, making all other factors extraneous.

The approach presented in this dissertation was designed to address the problem of tracking and identifying multiple people within their homes. It has shown that despite substantial measurement and environmental noise, the use of low-power wireless ultrasound sensors that measure height to do so is not only possible, but performs as well or better than existing technologies that require wearable devices, video cameras, or both.

The choice of height as a weak biometric was made because, while it is not a uniquely identifying feature worldwide, enough variance usually exists within a home to make it useful. By using height, the need for devices to be carried by participants is eliminated, as is the possibility for missing or incomplete data because the devices have been forgotten somewhere within the house. The chance

for privacy invasion due to video cameras is also eliminated, making it easier to convince people to agree to an installation of the system in their homes.

Indoor tracking systems provide the basis for many future technologies and smart-home applications, including elderly and patient monitoring, activity recognition, and micro-zoned heating and cooling. An elderly person could maintain independence within their own home without unduly worrying their descendants. Better care can be given to those suffering from depression or dementia through home health monitoring. A house can reduce energy consumption by targeting the specific rooms used and activities performed by occupants. All these applications and more may be possible with the further development of a system that tracks and identifies occupants of a house without requiring their active participation; a system that disappears into the woodwork.

## 7.1 Future Work

A number of parameters of the presented tracking system can be improved. A hardware re-design is necessary due to the restrictions of the current acoustic solution. Improvements to the signal processing algorithm will help it to better extract height and direction data from event streams.

The most obvious improvement to the hardware is the devising of a more precise sensor with a higher sample rate. An optical sensor may be the best solution, but creating one that will cover the entire doorway at the necessary resolutions is a challenge. Augmenting the height sensors with others to capture more weak biometric data, such as weight, would be another option. Additional information should aid the tracking system.

A pair of directional passive infrared sensors indicates motion on a particular side of the door. Incorporating this information in the weighting function could aid in directionality. A latch sensor that indicates whether a door is open or closed provides more information that could also be incorporated in the weighting function. This has the possibility to prevent incorrect doorway transitions, as well as eliminating noise. Both of these types of sensors could also aid in the energy reduction necessary to move towards a battery-powered solution.

An enabling technology, indoor tracking systems pave the way for the building of future systems reliant on the localization of house occupants. Smart zoning systems and activity recognition could be created to reduce a home's energy consumption. Monitoring systems could improve the quality of life for many. All of these applications are a possibility with continuing research of indoor tracking systems. The system presented in this dissertation still needs improvements, but it is a first step



towards realizing the above capabilities and many more that have not even been envisioned yet, all capable of improving the life of mankind.

# Bibliography

- [1] B.W. Parkinson, J.J. Spilker, P. Axelrad, and P. Enge. Global positioning system: Theory and applications volume ii. *Progress in astronautics and aeronautics*, 163, 1996.
- [2] American time use survey 2010. <http://www.bls.gov/news.release/archives/atus.06222011.pdf>, 2011.
- [3] D. Zhang, F. Xia, Z. Yang, L. Yao, and W. Zhao. Localization technologies for indoor human tracking. *Arxiv preprint arXiv:1003.1833*, 2010.
- [4] H. Koyuncu and S.H. Yang. A survey of indoor positioning and object locating systems. *IJCSNS*, 10(5):121, 2010.
- [5] L.M. Ni, Y. Liu, Y.C. Lau, and A.P. Patil. Landmarc: indoor location sensing using active rfid. *Wireless Networks*, 10(6):701–710, 2004.
- [6] K. Chawla, G. Robins, and L. Zhang. Efficient rfid-based mobile object localization. In *Wireless and Mobile Computing, Networking and Communications (WiMob), 2010 IEEE 6th International Conference on*, pages 683–690. IEEE, 2010.
- [7] R. Clarke. Human identification in information systems: Management challenges and public policy issues. *Information Technology & People*, 7(4):6–37, 1994.
- [8] C.H. Jiang and G.D. Su. Information fusion in face and fingerprint identity verification system. In *Machine Learning and Cybernetics, 2004. Proceedings of 2004 International Conference on*, volume 6, pages 3529–3535. IEEE, 2004.
- [9] Y. Zhu, T. Tan, and Y. Wang. Biometric personal identification based on iris patterns. In *Pattern Recognition, 2000. Proceedings. 15th International Conference on*, volume 2, pages 801–804. IEEE, 2000.
- [10] R.P. Wildes. Iris recognition: an emerging biometric technology. *Proceedings of the IEEE*, 85(9):1348–1363, 1997.
- [11] P.N. Belhumeur, J.P. Hespanha, and D.J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 19(7):711–720, 1997.
- [12] J. Jenkins and C. Ellis. Using ground reaction forces from gait analysis: body mass as a weak biometric. *Pervasive Computing*, pages 251–267, 2007.
- [13] W.H. Liao, C.L. Wu, and L.C. Fu. Inhabitants tracking system in a cluttered home environment via floor load sensors. 2008.
- [14] M.D. Addlesee, A. Jones, F. Livesey, and F. Samaria. The orl active floor. *IEEE Personal Communications*, 4:35–41, 1997.

- [15] R.J. Orr and G.D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *CHI'00 extended abstracts on Human factors in computing systems*, pages 275–276. ACM, 2000.
- [16] R. Headon and R. Curwen. Recognizing movements from the ground reaction force. In *Proceedings of the 2001 workshop on Perceptive user interfaces*, pages 1–8. ACM, 2001.
- [17] T. Yamazaki. Ubiquitous home: real-life testbed for home context-aware service. 2005.
- [18] Y. Kaddoura, J. King, and A.S. Helal. Cost-precision tradeoffs in unencumbered floor-based indoor location tracking. In *From smart homes to smart care: ICOST 2005, 3rd International Conference on Smart Homes and Health Telematics*, volume 15, page 75. Ios Pr Inc, 2005.
- [19] Z. Liu and S. Sarkar. Outdoor recognition at a distance by fusing gait and face. *Image and Vision Computing*, 25(6):817–832, 2007.
- [20] D. Gafurov and E. Snekenes. Gait recognition using wearable motion recording sensors. *EURASIP Journal on Advances in Signal Processing*, 2009:7, 2009.
- [21] SME Hossain and G. Chetty. Next generation identity verification based on face-gait biometrics. In *International Conference on Biomedical Engineering and Technology” Kuala Lumpur*, pages 17–19, 2011.
- [22] V. Srinivasan, J. Stankovic, and K. Whitehouse. Using height sensors for biometric identification in multi-resident homes. *Pervasive Computing*, pages 337–354, 2010.
- [23] R. Want and A. Hopper. Active badges and personal interactive computing objects. *Consumer Electronics, IEEE Transactions on*, 38(1):10–20, 1992.
- [24] Q. Hao, D.J. Brady, B.D. Guenther, J.B. Burchett, M. Shankar, and S. Feller. Human tracking with wireless distributed pyroelectric sensors. *Sensors Journal, IEEE*, 6(6):1683–1696, 2006.
- [25] D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 25(3-4):195–207, 1998.
- [26] L.E. Navarro-Serment, C.J.J. Paredis, P.K. Khosla, et al. A beacon system for the localization of distributed robotic teams. In *Proceedings of the International Conference on Field and Service Robotics*, volume 6. Citeseer, 1999.
- [27] P. Chen, P. Ahammad, C. Boyer, S.I. Huang, L. Lin, E. Lobaton, M. Meingast, S. Oh, S. Wang, P. Yan, et al. Citric: A low-bandwidth wireless camera network platform. In *Distributed Smart Cameras, 2008. ICDSC 2008. Second ACM/IEEE International Conference on*, pages 1–10. IEEE, 2008.
- [28] V. Menon, B. Jayaraman, and V. Govindaraju. Biometrics driven smart environments: abstract framework and evaluation. *Ubiquitous Intelligence and Computing*, pages 75–89, 2008.
- [29] K. Lorincz and M. Welsh. Motetrack: A robust, decentralized approach to rf-based location tracking. *Location-and Context-Awareness*, pages 49–62, 2005.
- [30] J. Schiff and K. Goldberg. Automated intruder tracking using particle filtering and a network of binary motion sensors. In *Automation Science and Engineering, 2006. CASE'06. IEEE International Conference on*, pages 580–587. IEEE, 2006.
- [31] Z. Wang, H. Li, X. Shen, X. Sun, and Z. Wang. Tracking and predicting moving targets in hierarchical sensor networks. In *Networking, Sensing and Control, 2008. ICNSC 2008. IEEE International Conference on*, pages 1169–1173. IEEE, 2008.
- [32] N.B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43. ACM, 2000.

- [33] A. Savvides and M.B. Srivastava. A distributed computation platform for wireless embedded sensing. In *Computer Design: VLSI in Computers and Processors, 2002. Proceedings. 2002 IEEE International Conference on*, pages 220–225. IEEE, 2002.
- [34] A. Savvides, C.C. Han, and M.B. Srivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 166–179. ACM, 2001.
- [35] L. Girod, M. Lukac, V. Trifa, and D. Estrin. The design and implementation of a self-calibrating distributed acoustic sensing platform. In *Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 71–84. ACM, 2006.
- [36] Y.M. Kwon, K. Mechitov, S. Sundresh, W. Kim, and G. Agha. Resilient localization for sensor networks in outdoor environments. In *Distributed Computing Systems, 2005. ICDCS 2005. Proceedings. 25th IEEE International Conference on*, pages 643–652. IEEE, 2005.
- [37] Y. Fukuju, M. Minami, H. Morikawa, and T. Aoyama. Dolphin: an autonomous indoor positioning system in ubiquitous computing environment. In *IEEE Workshop on Software Technologies for Future Embedded Systems*, volume 53, pages 53–56, 2003.
- [38] M. Minami, Y. Fukuju, K. Hirasawa, S. Yokoyama, M. Mizumachi, H. Morikawa, and T. Aoyama. Dolphin: a practical approach for implementing a fully distributed indoor ultrasonic positioning system. *UbiComp 2004: Ubiquitous Computing*, pages 347–365, 2004.
- [39] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. *Wireless Networks*, 8(2):187–197, 2002.
- [40] C. Taylor, A. Rahimi, J. Bachrach, H. Shrobe, and A. Grue. Simultaneous localization, calibration, and tracking in an ad hoc sensor network. In *Proceedings of the 5th international conference on Information processing in sensor networks*, pages 27–33. ACM, 2006.
- [41] C. Park and P.H. Chou. ecam: ultra compact, high data-rate wireless sensor node with a miniature camera. In *Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 359–360. ACM, 2006.
- [42] W. Feng, E. Kaiser, W.C. Feng, and M.L. Baillif. Panoptes: scalable low-power video sensor networking technologies. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCCAP)*, 1(2):151–167, 2005.
- [43] M. Rahimi, R. Baer, O.I. Iroezi, J.C. Garcia, J. Warrior, D. Estrin, and M. Srivastava. Cyclops: in situ image sensing and interpretation in wireless sensor networks. In *Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 192–204. ACM, 2005.
- [44] I. Downes, L.B. Rad, and H. Aghajan. Development of a mote for wireless image sensor networks. *Proc. of COGNITIVE systems with Interactive Sensors (COGIS), Paris, France*, 2006.
- [45] T. Teixeira, D. Lymberopoulos, E. Culurciello, Y. Aloimonos, and A. Savvides. A lightweight camera sensor network operating on symbolic information. In *ACM SenSys Workshop on Distributed Smart Cameras*. Citeseer, 2006.
- [46] R.G. Brown and P.Y.C. Hwang. *Introduction to random signals and applied Kalman filtering*, volume 2. John Wiley & Sons, 1997.
- [47] J. Krumm, S. Harris, B. Meyers, B. Brumitt, M. Hale, and S. Shafer. Multi-camera multi-person tracking for easyliving. *vs*, page 3, 2000.
- [48] P. Bahl and V.N. Padmanabhan. Radar: An in-building rf-based user location and tracking system. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 775–784. Ieee, 2000.

- [49] J. Liang, J. Shao, Y. Xu, J. Tan, B.T. Davis, and P.L. Bergstrom. Sensor network localization in constrained 3-d spaces. In *Mechatronics and Automation, Proceedings of the 2006 IEEE International Conference on*, pages 49–54. IEEE, 2006.
- [50] Y.W. Lee, E. Stuntebeck, and OC Miller. Merit: mesh of rf sensors for indoor tracking. In *Sensor and Ad Hoc Communications and Networks, 2006. SECON'06. 2006 3rd Annual IEEE Communications Society on*, volume 2, pages 545–554. IEEE, 2006.
- [51] D. Harmer, A. Yarovoy, N. Schmidt, K. Witrisal, M. Russell, E. Frazer, T. Bauge, S. Ingram, A. Nezirovic, A. Lo, et al. An ultra-wide band indoor personnel tracking system for emergency situations (europcom). In *Radar Conference, 2008. EuRAD 2008. European*, pages 404–407. IEEE, 2008.
- [52] W. ur Rehman, E. de Lara, and S. Saroiu. Cilos: a cdma indoor localization system. In *Proc. UbiComp*, pages 104–113.
- [53] V. Otsason, A. Varshavsky, A. LaMarca, and E. De Lara. Accurate gsm indoor localization. *UbiComp 2005: Ubiquitous Computing*, pages 141–158, 2005.
- [54] M. Azizyan, I. Constandache, and R. Roy Choudhury. Surroundsense: mobile phone localization via ambience fingerprinting. In *Proceedings of the 15th annual international conference on Mobile computing and networking*, pages 261–272. ACM, 2009.
- [55] National health statistics reports. <http://www.cdc.gov/nchs/data/nhsr/nhsr010.pdf>, 2008.
- [56] Improving pedestrian safety at unsignalized crossings. [http://onlinepubs.trb.org/onlinepubs/nchrp/nchrp\\_rpt\\_562.pdf](http://onlinepubs.trb.org/onlinepubs/nchrp/nchrp_rpt_562.pdf), 2006.
- [57] T.W. Hnat, V. Srinivasan, J. Lu, T.I. Sookoor, R. Dawson, J. Stankovic, and K. Whitehouse. The hitchhikers guide to successful residential sensing deployments. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2011.
- [58] 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.
- [59] Energy Information Administration. Table HC2-1a, Household Characteristics by Climate Zone, Million U.S. Households. [http://www.eia.doe.gov/emeu/recs/recs2001/hc\\_pdf/hous-char/hc2-1a\\_climate2001.pdf](http://www.eia.doe.gov/emeu/recs/recs2001/hc_pdf/hous-char/hc2-1a_climate2001.pdf), 2001. Accessed Sunday October 25, 2009.
- [60] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of the 4th international symposium on Information processing in sensor networks*, IPSN '05, Piscataway, NJ, USA, 2005. IEEE Press.
- [61] Maxbotix. <http://www.maxbotix.com/products/MB1200.htm>, 2011.
- [62] Synapse wireless: Snappy rf100. <http://www.synapse-wireless.com/>, 2011.
- [63] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, et al. Tinyos: An operating system for sensor networks. *Ambient intelligence*, 35, 2005.
- [64] Ping))). <http://www.parallax.com/>, 2011.
- [65] Passive infrared sensor. <http://www.parallax.com/>, 2011.
- [66] Go!motion. <http://www.vernier.com/products/>, 2011.

- [67] 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.
- [68] 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.
- [69] 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.
- [70] 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.
- [71] 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, 2003.(PerCom 2003). Proceedings of the First IEEE International Conference on*, pages 521–524. IEEE, 2003.
- [72] 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.
- [73] 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.
- [74] DJ Cook and M. Schmitter-Edgecombe. Assessing the Quality of Activities in a Smart Environment. *Methods of information in medicine*, 48(5):480, 2009.
- [75] 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.
- [76] W. Edwards and R. Grinter. At home with ubiquitous computing: Seven challenges. In *UbiComp 2001: Ubiquitous Computing*, pages 256–272. Springer, 2001.
- [77] S. Davidoff, M. Lee, C. Yiu, J. Zimmerman, and A. Dey. Principles of smart home control. *UbiComp 2006: Ubiquitous Computing*, pages 19–34, 2006.
- [78] 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.
- [79] 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.
- [80] 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.
- [81] 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.
- [82] J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. *Pervasive Computing*, 2006.