Improving RF Localization Through Measurement and Manipulation of the Channel Impulse Response

by

Benjamin Kempke

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Computer Science and Engineering)
in the University of Michigan
2017

Doctoral Committee:

Associate Professor Prabal Dutta, Chair
Professor David Blaauw
Associate Professor James W. Cutler
Professor Kang G. Shin

Benjamin Kempke

bpkempke@umich.edu

ORCID: 0000-0002-7908-183X

To my wife, my family, and my friends whose continual encouragement and support have inspired me to pursue and complete this work.

Acknowledgments

A gratifying and successful journey through the PhD program is not possible without a guide who is well-versed in the world of academia, adept in the art of research, and compassionate for any obstacles experienced along the way. To this end, my advisor, Professor Prabal Dutta has provided invaluable support and guidance towards identifying, executing, and completing research across a wide range of subjects throughout the last six years. Additionally, I thank my committee members, Professors James Cutler, Kang Shin, and David Blaauw, who comprise a diverse team of experts that have greatly helped shape the direction and organization of my dissertation.

Pat Pannuto, my colleague and right-hand man in brainstorming, evaluating, and articulating the localization concepts detailed in this dissertation has been paramount in its completion. Time and again he has provided candid advisement, supported late-night measurement campaigns, crafted captivating introductions, and seen many papers through their last-hour effort. Collectively, Lab11 has been an engaging and supportive group that truly understands the balance of work and play. Each individual in Lab11 has been selfless in providing assistance and promoting teamwork in ensuring the success of their coworkers.

Lastly, I thank my wife, my family, and my friends who have provided me with continual encouragement throughout my graduate years. The consideration from those close to me has been indispensable in maintaining a clear perspective and seeing tasks to completion. This has been a journey which would not have been possible without their support.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Improving RF Localization Through Measurement and Manipulation of the Channel Impulse Response

by

Benjamin Kempke

Chair: Prabal Dutta

For over twenty years, global navigation satellite systems like GPS have provided an invaluable navigation, tracking, and time synchronization service that is used by people, wildlife, and machinery. Unfortunately, the coverage and accuracy of GPS is diminished or lost when brought indoors since GPS signals experience attenuation and distortion after passing through and reflecting off of building materials. This disparity in coverage coupled with growing demands for indoor positioning, navigation, and tracking has led to a plethora of research in localization technologies. To date, however, no single system has emerged as a clear solution to the indoor localization and navigation problem because the myriad of potential applications have widely varying performance requirements and design constraints that no system satisfies. Fortunately, recently-introduced commercial ultra-wideband RF hardware offers excellent ranging accuracy in difficult indoor settings, but these systems lack the robustness and simplicity needed for many indoor applications. We claim that an asymmetric design that separates transmit and receive functions can enable many of the envisioned applications not currently realizable with an integrated design. This separation of functionality allows for a flexible architecture which is more robust to the in-band interference and heavy multipath commonly found in indoor environments.

In this dissertation, we explore the size, weight, accuracy, and power requirements imposed on tracked objects (tags) for three broadly representative applications and propose the design of fixed-location infrastructure (anchors) that accurately and robustly estimate a tag's location, while minimizing deployment complexity and adhering to a unified system architecture. Enabled applications range from 3D tracking of small, fast-moving micro-quadrotors to 2D personal navigation across indoor maps to tracking objects that remain stationary for long periods of time with near-zero energy cost. Each application requires careful measurement of the ultra-wideband channel impulse response, and an augmented narrowband receiver is proposed to perform these measurements. The key design principle is to offload implementation complexity to static infrastructure where an increase in cost and complexity can be more easily absorbed and amortized. Finally, with an eye towards the future, we explore how the increasingly crowded RF spectrum impacts current ultra-wideband system design, and propose an alternative architecture that enables improved coexistence of narrowband and ultra-wideband transmissions.

CHAPTER 1

Introduction

The advent of global navigation satellite systems (GNSS) provided a giant leap forward in providing accurate tracking, navigation and time synchronization in many differing application domains. GPS technologies are built into smartphones and cars to aid in navigating roads, making paper maps obsolete. Similar efficiencies have been brought to the tracking and navigation of satellites, airplanes, and seafaring vessels. In addition to its uses as a navigational aid, recent efforts have shown that automated control of these vehicles is possible, allowing for safer waterways, roadways, and skies by minimizing the dangers of inherent human error.

While GPS provides sufficient tracking capabilities in outdoor environments where few objects obstruct the sky, indoor environments pose a difficult arena in which GPS signals undergo significant attenuation and distortion from building materials, humans, equipment, furniture, and other objects present in the environment. As such, personal navigation through unfamiliar indoor environments usually requires one to fall back to posted maps. Alternatively, the automated tracking and navigation of indoor robots typically employs the use of expensive optical scanning technologies or error-prone dead reckoning techniques. For robotic navigation applications which are unable to afford the cost of expensive optical solutions, a careful consideration of the performance requirements is required. Furthermore, there exists an entire class of objects (assets) whose knowledge of location is important, yet are left entirely unserved due to stringent requirements on the size, weight, cost, and power imposed on any affixed hardware. Throughout this dissertation, we will explore a variety of system design concepts for addressing these and other indoor localization scenarios.

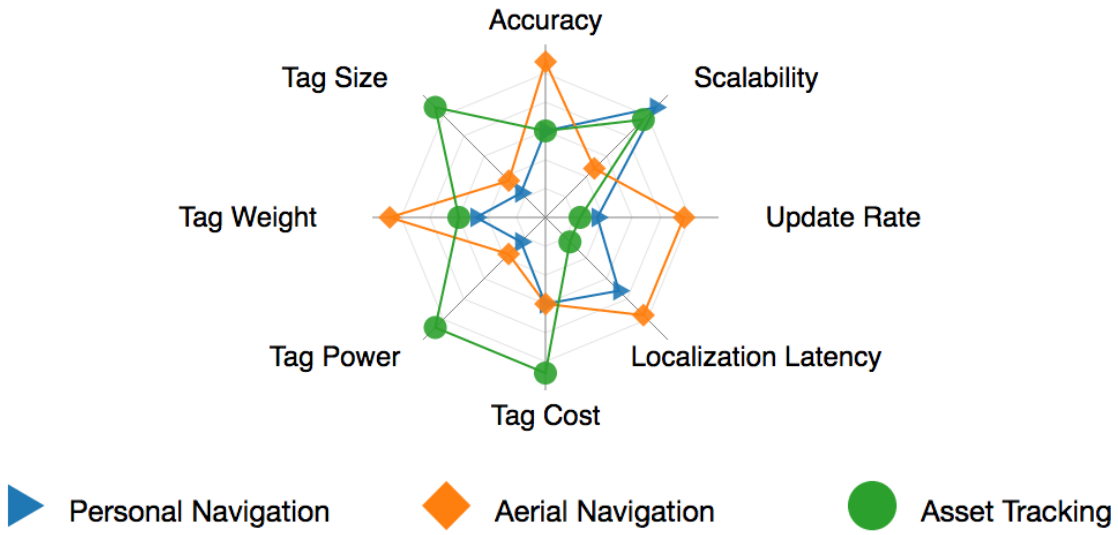

Figure 1.1: **Figures of Merit for Indoor Localization.** A qualitative consideration of the relative importance of common figures of merit in representative indoor localization applications, where larger values are considered to be more important. Due to the vast difference in priorities different tag designs are necessary to meet these requirements, and the impact of any associated infrastructure should be carefully considered in evaluating the impact on system deployment.

1.1 A Search Engine for the Physical World

The field of localization covers a wide array of services aimed at providing knowledge of an object's physical location. Similar to the internet search engine in the cyber domain, localization provides answers to physical queries such as "Where are my keys?" or logical questions such as "In what building am I?" With accurate knowledge of an object's physical location, higher-level services like assisted navigation are possible.

Common metrics used to evaluate localization performance are shown in Figure 1.1. Characteristics of the anchor (localization infrastructure) and tag (the physical object being tracked) must be considered in evaluating the importance of each metric. In order to evaluate the importance of each of these metrics in designing a system that achieves indoor localization, the set of applications to be supported must be considered.

1.2 Motivating Indoor Localization Applications

A number of applications for indoor navigation and tracking have been proposed that have the potential to greatly impact the way in which we live and interact with the indoor environment. Smartphone-based navigation of indoor environments would enable the navi-

gation of unfamiliar environments in much the same way that GPS has revolutionized our navigation of outdoor spaces.

In the same way that accurate indoor localization will enable a bevy of personal navigation applications, the potential effects of the accurate navigation of autonomous indoor robots are far-reaching. Automation of industrial warehouses through automated inventory tracking and transportation is a valuable localization use-case. Forward-thinking applications such as the aerial delivery of goods throughout the indoor environment, or the augmentation of floor-based indoor robotic systems, have the potential to aid in more and more facets of everyday living.

Furthermore, extremely low-power "energy harvesting" sensors have received considerable interest due to their ability to offer computing and sensing on battery-less nodes. These devices harvest energy from a variety of sources in the indoor environment including light, radio frequency, and vibration, among others. Their ability to harvest energy from their surrounding environment allows them to forego battery replacement and allows them to be placed throughout the indoor environment without being limited to locations where wired power is available. The potential for small, energy harvesting tags would allow us to query for the location of our lost keys at the touch of a button or query the location of inventory in a large warehouse.

Despite the numerous motivating applications of indoor localization, no solution for high-fidelity indoor localization that satisfies these diverse applications currently exists. The performance requirements for indoor localization vary greatly based on the application, and no system currently exists which is able to meet these widely varying demands. Therefore, a deeper investigation into the challenges of achieving these performance criteria is required.

1.3 Challenges of Indoor Localization

To date, no system has demonstrated adequate flexibility and performance to support the wide range of indoor localization applications that have been proposed. Those systems which are "infrastructure-free" (requiring the addition of no dedicated hardware) typically have at least one of the following drawbacks which hinder their broader adoption: a) they require a large training effort b) they are expensive, or c) they are very sensitive to small changes in the indoor environment. Due to the absence of a broadly usable, infrastructure-free solution, much effort has been focused on the adoption of infrastructure-based solutions that aim to balance the system's performance and applicability across the broadest range of potential applications. However, to date, the widescale adoption of infrastructure-

based solutions has also remained elusive, perhaps due to the high costs of deployment and retrofit of hardware necessary to support large scale deployment. We believe that to facilitate the widespread adoption of an indoor localization retrofit, the hardware costs must be low, the system must require minimal setup effort and maintenance, and the architecture must be flexible enough to support the most common applications of indoor localization.

1.3.1 Indoor Localization Infrastructure Challenges

Major infrastructure-based indoor localization technologies can be divided into three different categories based on the medium used for localization operations. The three media used for localization are sound, light, and radio frequency (RF).[1] Systems that perform localization using light and sound are at a disadvantage compared to RF localization schemes due to the prevalence of opaque objects in the indoor environment. However, issues related to the inaccuracy of RF localization systems has led to the use of sound and light in applications which require centimeter-level accuracy.

Historically, RF localization methods have been plagued by poor accuracy due to a mismatch between the time-domain resolution of commodity radio signals to the speed-of-light propagation of radio signals in air. Angle-of-arrival techniques such as those employed by optical methods require complex and large beamforming elements which greatly increase the system's deployment cost and complexity. Received signal strength (RSSI) techniques suffer from large errors due to the effects of interference in multipath environments. Time-domain techniques have typically shown the best performance in static settings, but their accuracy is limited due to the coarse time-domain resolution afforded by the narrow bandwidths utilized by off-the-shelf radios.

The recent emergence of ultra-wideband[2] standards and legislation along with low-cost, standards-compliant hardware has opened up new opportunities to realize accurate RF time-based localization techniques. Advertised ranging accuracy of these radios is around 10 cm, potentially enabling the realization of localization systems that are competitive with the accuracy and performance afforded by sound- and light-based techniques, without many of the drawbacks of those systems.

[1] Additional measurements including magnetics or vibration can also be found in indoor localization research, yet they are typically used to augment sound, light, or RF measurements, and are thus considered irrelevant for the sake of this discussion.

[2] Radios are considered "ultra-wideband" if their transmissions exceed the lesser of 500 MHz absolute bandwidth or 20% fractional bandwidth

1.3.2 Indoor Localization Tag Challenges

Localization applications bring up a wide range of competing design points when considering metrics such as performance, cost, size, and weight of the burden applied to the object being tracked. In order to support the widest range of applications, we believe that as much of the burden as possible should be shifted to fixed infrastructure to simplify the localization tag. At the same time, the infrastructure financial and deployment burden must remain viable to not prohibit widespread adoption and general enough to support a sufficiently diverse array of localization applications to justify the costs.

Allowable tag power may vary by many orders of magnitude. Similar to wireless sensor networking applications, the tags may choose to utilize high-power, active transceivers (drawing *milliwatts*) for fast, fine-grained spatial accuracy, or trade this off for slower update rates with more constrained backscatter links (drawing *microwatts*). The underlying system must adopt a flexible localization protocol which is able to support these wide-ranging tag power requirements, perhaps with attendant trade-offs in accuracy or latency.

In a similar manner, metrics such as tag cost, size, and weight can provide a rich tradeoff space for the tag's design. For instance, a consumer-centric asset localization application may require small, inexpensive tags whereas an industrial asset localization application may be able to sacrifice node size and weight for higher localization performance. In this work, we will propose a variety of different tag designs which aim to cover a a diverse set of indoor localization requirements.

1.4 The CIR and the Bandwidth-Accuracy Tradeoff

The most accurate RF-based localization techniques [29] utilize large amounts of bandwidth and attempt to determine position through an accurate determination of the arrival time of exchanged radio messages. Ultra-wideband techniques achieve the highest accuracy in determining time-of-arrival due to their ability to accurately distinguish individual paths of a multipath propagation that is typical in indoor environments.

Accuracy determines the resolution at which the channel impulse response can be recreated. As such, ultra-wideband techniques have shown the highest accuracy. However, with a typical UWB channel bandwidth of 500 MHz, many different paths of multipath propagation can be present within the time-domain resolution of the UWB channel's channel impulse response.

Figure 1.2 shows the time-domain representation of a typical indoor RF multipath channel. By imposing a bandwidth restriction on the channel, many of the first paths of mul-

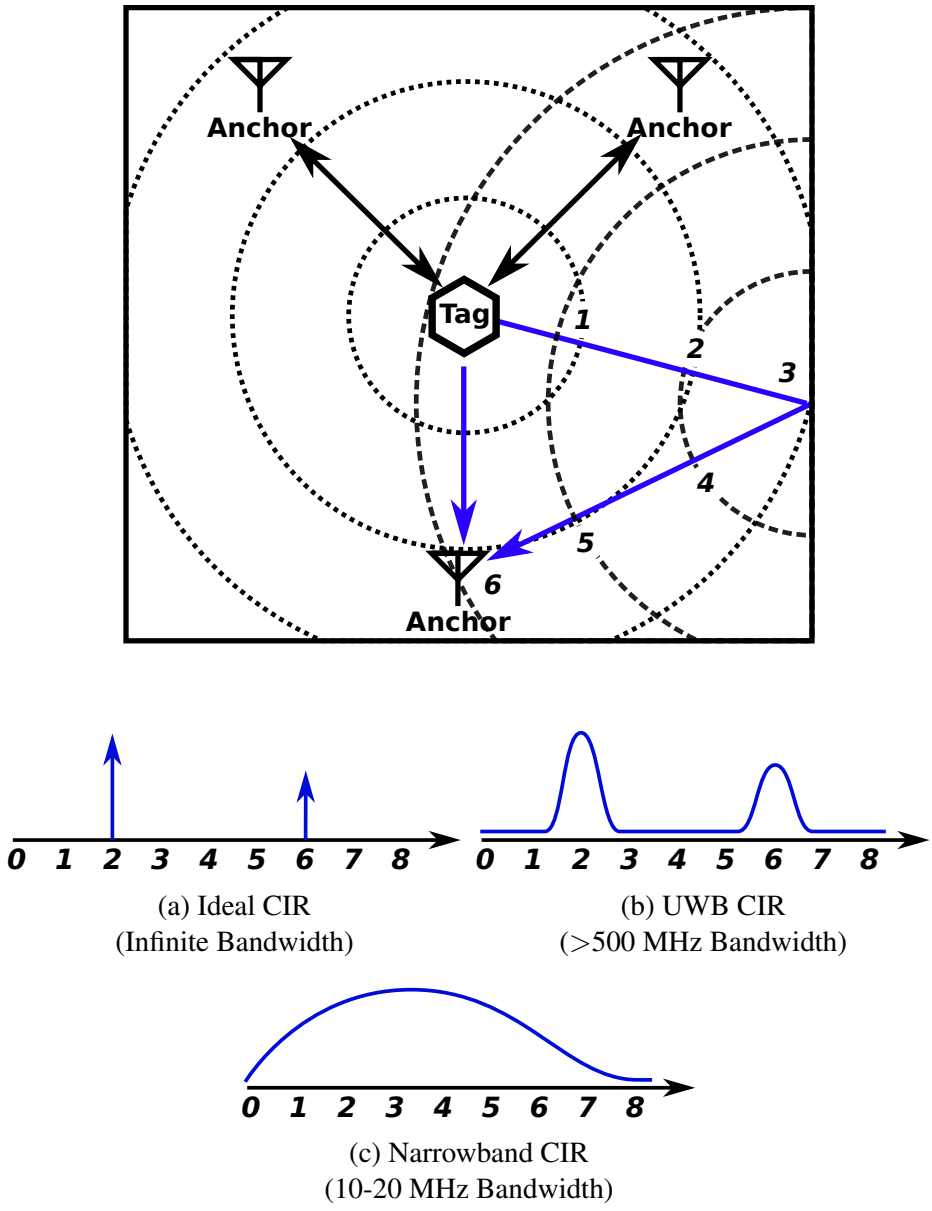

(a) Ideal CIR
(Infinite Bandwidth)

(b) UWB CIR
(>500 MHz Bandwidth)

(c) Narrowband CIR
(10-20 MHz Bandwidth)

Figure 1.2: **The Indoor Channel Impulse Response.** Graphical depiction and resulting channel impulse response for a representative indoor multipath environment observed between a tag and fixed-location infrastructure. A tradeoff occurs between the distinguishable time-domain accuracy of the line-of-sight path and bandwidth, yielding a direct correlation to the high accuracy achieved by ultra-wideband localization technologies in complex indoor multipath environments.

tipath propagation coalesce, and therefore distort, the bandwidth-limited representation of the RF channel. The amount and magnitude of this distortion is determined by the magnitude, phase, and distribution of coalescing multipath on top of the line-of-sight path.

1.5 Suitability of Ultra-Wideband Technologies

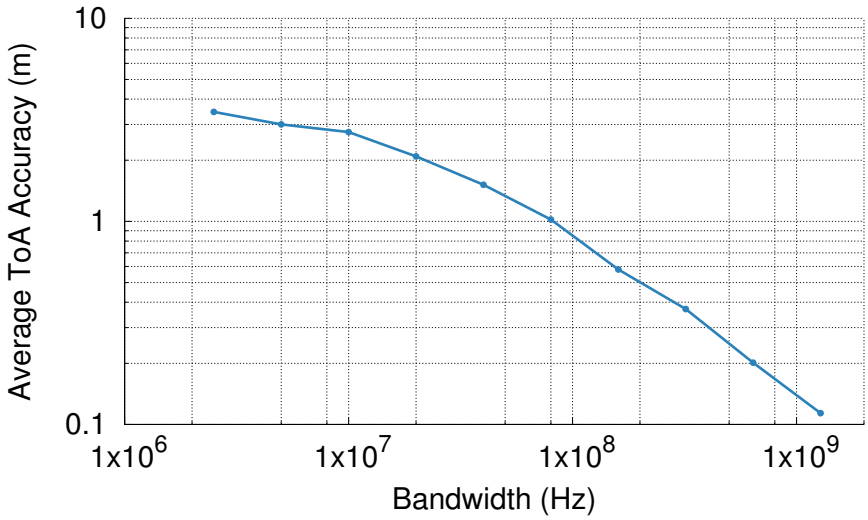

Figure 1.3: **Time-of-Arrival Accuracy vs. Bandwidth.** The heavy multipath experienced in indoor RF environments causes difficulty in determining the time-of-arrival for RF signals, an important metric for characterizing indoor localization performance. Shown is the average ranging accuracy achievable with the leading line-of-sight time-of-arrival determination algorithm using the widely-referenced 802.15.4a channel model [62]. Time-of-arrival accuracy realizes an approximately linear relationship to the amount of utilized bandwidth past 80 MHz. This is due to the separation of interfering multipath as resolution increases from a corresponding increase in bandwidth. This has led to much interest in utilizing large amounts of bandwidth (so-called "ultra-wideband" channels) to maximize localization accuracy in indoor environments.

RF-based localization technologies offer the best coverage and cost for indoor environments, but they have historically shown poor accuracy in these settings due to difficult multipath characteristics. The room-level accuracy afforded by narrowband RF localization techniques has not been sufficient for fine-grained localization applications, leading to the utilization of large amounts of RF bandwidth. Figure 1.3 shows the tradeoff between localization accuracy and utilized bandwidth. Recent developments in commercial ultra-wideband transceivers have made the sub-meter accuracy required by most indoor localization applications possible. However, little attention has been placed on the tag cost, complexity, and power for those applications in which these performance metrics are important.

1.6 Thesis Statement

Careful manipulation and segmented measurement of the ultra-wideband channel from different frequencies and vantage points permits flexible and deployable localization systems with large improvements in achievable tag size, weight, accuracy, and power due to the large bandwidth and time-domain sparsity of ultra-wideband transmissions along with the spatial and polarization diversity afforded by multiple, inexpensive antennas.

1.7 Contributions of this Dissertation

We present the design, implementation, and evaluation of a number of architectures that provide fast and accurate indoor localization with widely varying tag performance capabilities, enabling a number of useful indoor localization applications. The remainder of this dissertation will describe the concepts, implementation, and evaluation of the capabilities described in the thesis statement.

In Chapter 2, we begin identifying a number of diverse, motivating applications for indoor localization, along with the various constraints (accuracy, power, size, weight, update rate) they impose on the underlying localization system. From this discussion, a number of performance criteria are identified for each application. These performance criteria, along with those relating to adoption (system cost, deployability, etc.) are used to evaluate the suitability of prior RF localization system designs and identify the drawbacks with these earlier approaches.

Chapter 3 identifies the hardware design choices underlying our RF localization system called SurePoint. As the most provisioned localization system architecture described in this dissertation, each SurePoint node uses a common hardware module with dedicated radio and microcontroller to hide complex implementation details and provide a simple I^2C interface that exposes raw ranging estimates to the host system. To aid in achieving high accuracy in indoor environments, SurePoint was designed to operate using commercially-available ultra-wideband transceivers. To minimize the occurrence of large ranging errors inherent to multipath-rich environments, three antennas are placed on each node to provide the option for many, independent measurements. This system and its underlying concepts were presented at HotWireless'15 [44] and SenSys'16 [42] and is joint work with Pat Pannuto and Prabal Dutta.

In Chapter 4, we identify a number of situations and applications in which the original hardware design would fail to meet the specified requirements of one of the applications detailed in Chapter 2. For these situations, we detail modifications to the hardware that

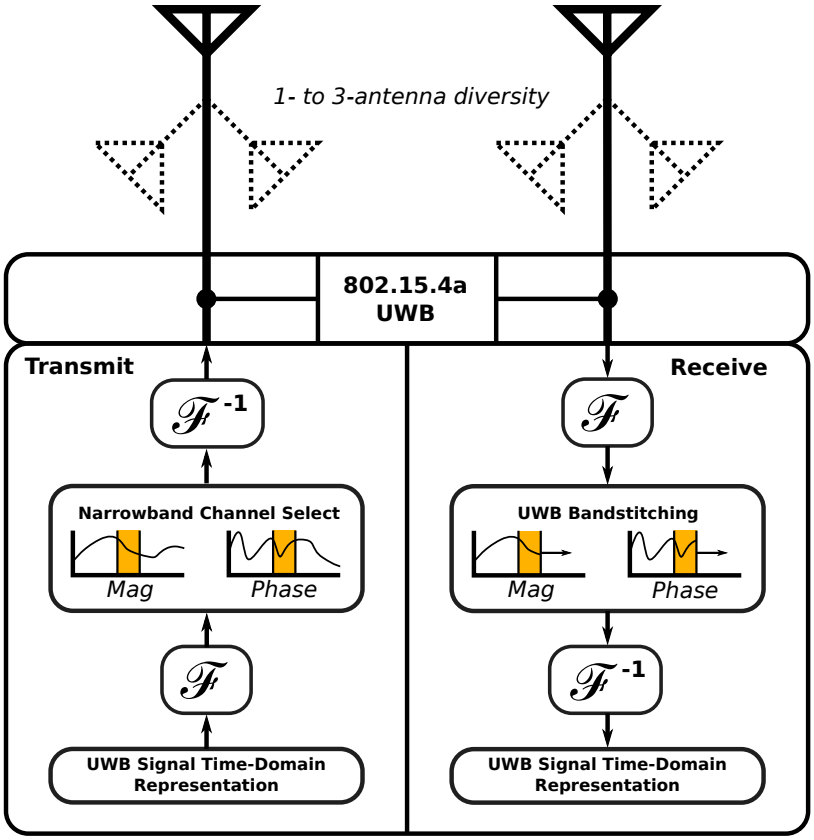

(a) Anchor Architecture

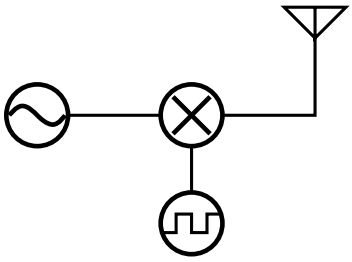

(b) Retrofit UWB Tag

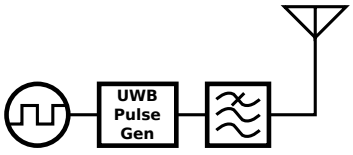

(c) Direct UWB Tag

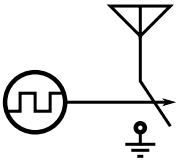

(d) Backscatter Tag

Figure 1.4: **Contributions of this Dissertation.** Ultra-wideband techniques and regulations have laid the groundwork for indoor localization with high coverage and accuracy. Recent standards have given rise to commercially-available ultra-wideband-compliant transceivers. However, a number of implementation challenges occur with many motivating applications of indoor localization. The concepts covered within this dissertation are shown here, ranging from the detailed analysis of the ultra-wideband channel impulse response, low-impact methods for its measurement, and various methods for achieving localization given wide-ranging tag size, weight, accuracy, and power requirements.

would aid in solving these issues. Systems designed with these modifications were presented at HotWireless'14 [43] and IPSN'16 [45] and is joint work with Pat Pannuto and

Prabal Dutta. Major issues covered include the design of low-cost, transmit-only tags for fast object tracking along with potential modifications to the anchor design to make them less susceptible to narrowband interference.

In Chapter 5, we introduce an energy-harvesting tag design that utilizes ultra-wideband backscatter modulation to achieve sub-microwatt, batteryless operation. By trading off update rate for tag power, the system offers a solution for obtaining accurate location estimates of rarely-moving objects. Furthermore, a new backscatter receiver design is presented which is easily implemented on traditional software-defined radios. This work, a collaboration with Pat Pannuto and Prabal Dutta, has recently been submitted for peer review.

Finally, in Chapter 6, we identify a number of drawbacks of traditional UWB infrastructure requirements and identify ways in which the drawbacks can be addressed. Specifically, the concept of UWB flooding is introduced, along with a description of its applicability in enabling MAC-layer primitives like global time synchronization and multi-tag schedule coordination. Additionally, its possible use in providing nanosecond-level time synchronization is proposed, and a variety of potential use cases for this new primitive is explored. These deployment aides were applied to the SurePoint system and presented at SenSys'16.

CHAPTER 2

Background and Related Work

Localization – the process of determining an object's absolute, relative, or semantic position has experienced a varied past incorporating many different technologies across many diverse applications without one clear, universal solution. GNSS technologies have provided absolute, global positioning capabilities in outdoor environments, enabling car-level navigation, but fall short in providing fast, sub-meter estimates and fail completely in most indoor environments. For this reason, a myriad of localization technologies and implementations have been developed to bridge the gap in situations which have demands unfulfilled by the GNSS approach.

A number of potential applications for accurate localization have been identified in previous work [15, 56, 61, 70] for both indoor and outdoor settings. Those which are viewed to be particularly useful are the applications encompassing person navigation through indoor environments (similar to outdoor vehicular navigation), robotic navigation, and asset tracking. A number of technologies have been proposed for the purposes of enabling these applications, and the following sections analyze the applicability of each of these proposed solutions.

2.1 Quadrotors, Ourselves, and Stuff: Motivating Indoor Localization

A number of interesting applications which help to drive the development of localization technologies in environments where GNSS solutions *are not feasible* are discussed in the following sections, along with the rough, order-of-magnitude performance requirements they impose on the underlying localization technology. While these constraints may not be true for every end-user application, they provide a rough idea of the constraints imposed on the underlying localization technologies they are based on. A summary of these results is given in Table 2.1.

(a) Personal Navigation

(b) Robotic Navigation

(c) Asset Tracking

Figure 2.1: **Quadrotors, Ourselves, and Stuff: Useful Applications for Indoor Localization.** Personal navigation – or the guidance through indoor environments – can reduce the time required to navigate unfamiliar buildings. Similarly, robotic navigation can allow for better autonomous navigation of both land and air-based robots through indoor spaces. Asset tracking – or the localization of one's stuff – can be used to track inventory in large industrial warehouses or in the residential setting as an aid in finding lost property. No one system has yet been proposed which is able to support these diverse applications while maintaining acceptable impact on size, weight, accuracy, and power of any affixed tags.

Personal Navigation. Personal navigation includes the aided navigation of indoor and outdoor environments with turn-by-turn directions to reach an intended destination on foot. Although the problem of outdoor navigation has been primarily solved with GNSS systems, the problem of personal navigation in indoor environments has yet to find a suitable solution.

Personal navigation requires absolute positioning capabilities with an approximately 1 m accuracy and 1 Hz update rate in order to easily navigate indoor spaces to doorway-level accuracy. With an intended end-user augmentation of the conventional smart phone, tag power and size constraints can be fairly relaxed, at around 1 W of active power and 5 cm^3 of standalone volume.

Robotic Navigation. Robotic navigation applications include the autonomous navigation of a variety of robotic vehicles throughout the indoor environment, including such platforms as the Roomba vacuum cleaner in residential settings, remote presence robots such as the RP-VITA, or the navigation of aerial vehicles. Unfortunately, the size, weight,

power, and accuracy requirements for robotic navigation can vary widely, so a select few case studies have been selected for analysis in this section.

Under consideration are the applications of flying micro-quadrotors and land-based robotic navigation. These have more stringent accuracy requirements (approximately 10 cm) than personal navigation, while update rate requirements can vary widely. In the case of land-based robotic navigation, order-Hz update rates are more than sufficient, whereas flying objects need order-10 Hz update rate in order to carefully navigate throughout the indoor environment without concern. While again application-dependent, tag power and size should be kept below 100 mW and 5 cm^3, respectively.

Asset Tracking. Real-time asset tracking is valuable in a wide range of applications including industrial manufacturing plants and warehouses, along with finding valuable or often-misplaced personal effects, to name a few. The performance criteria for these applications can vary greatly from those applications discussed so far, as most assets are typically stationary for significant periods of time. This decreases the required position update rate, but decreases the per-unit price and power target for the tracking of (potentially many) assets.

2.2 Performance Criteria for Localization Technology

Table 2.1 lists rough, order-of-magnitude performance criteria requirements for the applications mentioned in the last section. These performance criteria will change on a case-by-case basis, but these requirements will guide the design and selection of enabling technologies in these widely-varying application domains.

For the purposes of discussion in this dissertation, we will focus on the major issues of tag power, localization accuracy, and tag integrability for each of the identified localization application domains.

Performance Metric	Personal Navigation	Robotic Navigation	Asset Tracking
Accuracy [70]	1 m	10 cm	1 m
Update Rate	1 Hz	10 Hz	0.001 Hz
Tag Power	1 W	100 mW	1 μW
Tag Size	10 cm^3	5 cm^3	1 cm^3

Table 2.1: **Performance Criteria for Personal Navigation, Robotic Navigation, and Asset Tracking Applications.** These wide-ranging yet useful applications impose stringent criteria on specific localization performance metrics. In order to aid in its adoption, any deployed hardware must be able to support the operations and analyses required to attain the level of performance presented here. Estimates of update rate, tag power, and size take into account the relative mobility and burden allowable for the constrained use-cases identified in this dissertation.

2.3 Suitability of Current Localization Technologies

Localization technologies are a well-studied topic with immeasurable research spanning many different domains. The most well-studied domains include that of optical, acoustic, narrowband RF and ultra-wideband RF techniques. While localization using different media including magnetics and inertial measurements have been studied, their applicability has been severely limited due to their inaccuracies and susceptibility to change in dynamic environments.

The following sections detail the strengths and weaknesses of localization utilizing each medium, along with their applicability to each of the application domains described above.

2.3.1 Optical Localization

Optical localization technologies utilize light (or near-field infrared) to locate an object with recognizable optical characteristics. Passive techniques such as those used by motion capture systems utilize retroreflective markers which are identified by cameras located throughout the environment [65, 84]. At least two cameras are necessary to track an object in three dimensions with localization operations being based on triangulation. The accuracy of motion capture systems is typically less than one millimeter, and are used extensively and primarily in the film industry due to their high accuracy and responsiveness but high system cost.

Situations in which the tag is active give rise to a different class of optical localization technologies. Examples of active optical localization tags include those which track based on emission of light [48], or those which utilize a camera to identify optical ‘anchors’ with

(a) Optical Motion Capture

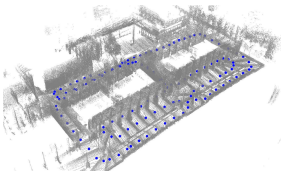

(b) Visual SLAM

(c) Optical Beacon

Figure 2.2: **Optical Localization Methodologies.** Many localization technologies have been developed which utilize light-based sensing to determine an object's position. Optical motion capture tracks reflective markers using stationary cameras. Optical beacon techniques determine position by tracking markers placed throughout the environment with a camera affixed to the object being tracked. VSLAM techniques determine position through an analysis of the optical characteristics of the environment itself. Optical techniques are often expensive or high-power, limiting their applicability to many consumer-centric applications.

known locations posted throughout the environment [68].

Optical localization technologies meet many of the performance criteria provided in the last section, but have a few distinctive disadvantages compared to other localization methodologies. First, optical techniques require line-of-sight between the tag and surrounding environment. Any occlusions of the tag being tracked result in the inability to update the tag's location. Second, those systems which support passive tags (i.e. low-to-no power) are typically unable to uniquely identify the tag being localized. If the passive tag is to go out of sight, the newly-acquired tag cannot be uniquely identified. Lastly, optical localization technologies are expensive, requiring the use of calibrated cameras to perform localization operations.

2.3.2 Narrowband RF Localization

RF localization techniques (both narrowband and wideband) benefit from the ability of RF to propagate through many objects which are visually and acoustically opaque. This allows for greater coverage and less infrastructure requirements when compared to the acoustic and optical localization technologies covered so far.

Narrowband RF techniques also benefit from a ubiquitous presence in electronic devices today. Narrowband transceivers typically follow a common architecture, allowing for minimal (to no) modifications to current transceiver design in order to support the measurements necessary for RF localization tasks. A majority of the techniques used for narrowband localization can be broken up into three categories: those which are RSSI-based,

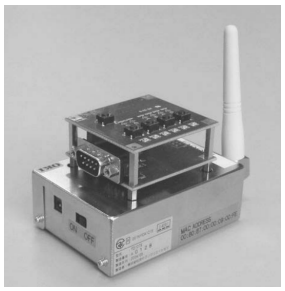

(a) RSSI
(with ZigBee [75])

(b) Fingerprinting
(Ekahau RTLS [40])

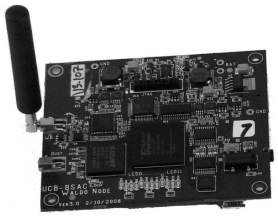

(c) Time-of-flight
(Waldo [50])

Figure 2.3: **Narrowband Localization Methodologies.** A large bevy of research has been performed on achieving accurate indoor localization utilizing existing narrowband architectures. Without the use of complex electronics (beamforming) or requiring movement to inform time-domain analysis, narrowband techniques do not have adequate time resolution to inform high accuracy measurements in multipath-rich environments.

fingerprint-based, and those which utilize time-of-flight.

Heavy multipath present in many indoor environments significantly affects the quality of localization estimates provided by narrowband approaches. This is due to the limited time-domain resolution of narrowband signals (2 meters or more). Time-based narrowband approaches are thus unable to reliably achieve better than 1 meter accuracy in indoor environments [93]. RSSI-based schemes provide worse performance based on their inability to separate the effects of multipath from that of the line-of-sight path. Fingerprinting-based schemes have shown better performance than their other narrowband counterparts, but require the use of large amounts of training data and are invariably prone to inaccuracies due to dynamic changes in the indoor environment.

Some systems have attempted to improve accuracy through the combination of multiple techniques [6] or the fusion of many sensing modalities [47] yet this tends to add significant complexity with minimal benefit. Advanced phase-based multi-antenna techniques [94, 96] have also been demonstrated to improve narrowband performance in indoor settings but are bulky and complex. Further systems have attempted to forego the estimation of geoposition in favor of semantic localization [13, 20], movement detection [1], or drawing in thin air [87], all alternative forms of localization not covered in this dissertation.

2.3.3 Acoustic/Ultrasonic Localization

Acoustic localization techniques utilize many of the same principles as their narrowband RF counterparts mentioned previously, but are able to take advantage of the much slower

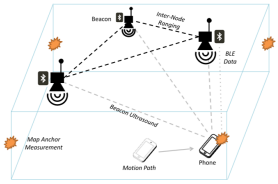

(a) Passive Ultrasonic (ALPS [51])

(b) Active Ultrasonic (Active Bat [88])

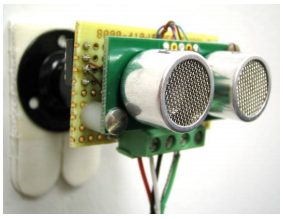

(c) Ultrasonic Echolocation [86]

Figure 2.4: **Acoustic Localization Methodologies.** Localization technologies utilizing acoustic techniques have been gaining popularity due to the prevalence of acoustic transducers in smartphones. Various techniques have been proposed including near-ultrasonic geopositioning with time-of-flight to echolocation approaches which calculate position based on the expected echo "fingerprint" of the surrounding environment. However, tag cost and the coverage afforded by these acoustic approaches have limited their applicability.

propagation speed of sound through air. With a propagation speed of approximately 340 m/s, a bandwidth of only 340 Hz is required to obtain a channel impulse response resolution of one meter, while many systems take advantage of *several kilohertz* of acoustic bandwidth while performing localization operations.

Figure 2.4 shows a selection of ultrasonic platforms for indoor localization which span the three techniques used for performing localization in the acoustic setting. Passive ultrasonic techniques [51, 66] employ acoustic receivers at the tag which calculate time difference measurements of transmissions coming from anchors situated throughout the environment. Active techniques [88] generate range estimates by employing a 2-way ranging protocol, adding the requirement of ultrasonic transmission at each tag. Finally, echolocation approaches [78] are anchor-less, and calculate position based on the expected acoustic channel impulse response of the surrounding environment.

Acoustic localization systems occupying the audible bands are often undesirable due to their impact on the perceived acoustic environment, so ultrasonic techniques have historically received the most attention. However, as with optical techniques, the acoustic environment is easily occluded, leading to decreased coverage and the need for many anchors in dense indoor environments. Also, acoustic localization schemes have not been able to achieve ultra low-power operation and are comparatively bulky, limiting their use in ultra-constrained application domains.

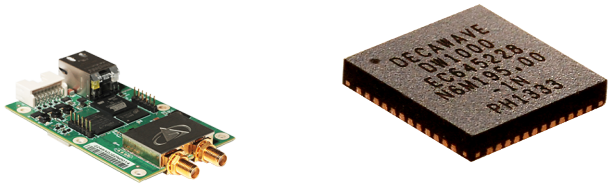

(a) Commercial UWB (Time Domain [21])

(b) DecaWave IC (DecaWave [18])

Figure 2.5: **Localization-Capable UWB Transceivers.** Commercial UWB localization platforms have been around for decades, yet only target the industrial domain due to the high cost of custom hardware required for their use. Recently, the DW1000 has been introduced which is the first off-the-shelf UWB transceiver targeting the localization domain, yet has limitations to its applicability and is too costly for high-volume localization applications.

2.3.4 Ultra-Wideband RF Localization

In order to bridge the gap between the spatial resolution of narrowband RF (>2 m) and that of acoustic techniques (<10 cm), further localization technologies have moved toward the adoption of large RF bandwidths (>500-1000+ MHz). This can yield spatial resolutions on the order of those achievable with acoustic techniques, but with the advantage of easier, unobtrusive adoption into current electronic systems and less obstructions due to "opaque" objects situated throughout the environment.

Regulations enacted over the last 15 years [25, 39] have enabled the use of ultra-wideband transmitters at low transmit power levels, enabling the realization of high fidelity localization in indoor environments without any significant effect on narrowband radios occupying the same spectrum. The adoption of ultra-wideband techniques has, however, been slow due to the high price target of current end-to-end localization systems and an inability to adequately handle the destructive effects of narrowband interference.

A variety of different techniques to utilize UWB bandwidths have been studied in the past [82, 90, 91]. Industrial applications have been using costly, custom-built UWB localization systems [21, 81] for many years. The custom-built and high-power nature of these systems has limited their general applicability beyond that of defense and industry. Other systems have improved localization accuracy by leveraging multiple ISM bands across large bandwidths [82]. Finally, the WASP system [71] did not achieve UWB-level bandwidths, but demonstrated the ability to perform bandstitching across successive narrowband bandwidths, a technique leveraged by Harmonium in Chapter 4.

2.3.5 Magnetic Localization

The magnetic medium has seen widespread interest in its ability to provide localization capabilities due to Earth's permeating geomagnetic field [46]. However, these approaches are only able to achieve room-level accuracy due to distortions in the magnetic field due to building materials and rely on large amounts of training effort to identify the building's characteristic footprint.

	Attenuation (dB)	
Material	**2.4 GHz**	**5.8 GHz**
Interior drywall	3-4	3-5
Cubicle wall	2-5	4-9
Wood door (hollow – solid	3-4	6-7
Brick/concrete wall	6-18	10-30
Glass/window (not tinted)	2-3	6-8
Double-pane coated glass	13	20
Bullet-proof glass	10	20
Steel/fire exit door	13-19	25-32

Table 2.2: **Coverage Aspects of RF Localization Technologies.** Many building materials commonly found in indoor environments are transparent to RF energy. For this reason, RF-based localization technologies achieve much higher coverage and infrastructure can be unobtrusively deployed (e.g. behind drywall) in comparison to other localization media such as light and sound [41].

2.4 UWB RF: Accurate Localization with High Coverage

Of the technologies discussed in the previous section, ultra-wideband localization systems provide the best application coverage with a reasonable per-unit cost and minimal infrastructure requirements. This is due to the wide coverage capabilities in indoor environments, unobtrusive integration into pre-existing electronic systems, and the high time-domain resolution of ultra-wideband signals. For this reason, UWB localization algorithms have relied on time-of-arrival measurement for localizing tags.

2.4.1 UWB Localization Algorithms

Depending on the problem formulation, a number of different techniques can be used to determine a tag's position given UWB transmit and/or receive timing information. Most techniques require the use of fixed-location infrastructure to frame position estimates in a specific coordinate frame.

Trilateration can be used to calculate a tag's position based on range estimates to nearby infrastructure (tags with receive and transmit functionality). Multilateration is used when only receive timing information is known at the anchors (transmit-only tags). Other time-difference techniques can be utilized for reflector-based designs (backscatter tags). These algorithms are well-known, and have been studied extensively in prior literature.

The remainder of this section will explore the technology space, identify shortcomings of the technology in its current state, and propose the design of ultra-wideband hardware which is able to best cover the applications proposed so far.

2.4.2 State-of-the-Art

Ultra-wideband localization has a storied past, spanning many different methods and commercial realizations which have been met with moderate success. An analysis of what has been performed in the past along with an identification of the pitfalls of the approaches taken so far will help to guide the design of systems with wider deployability and applicability in the future.

Commercial ultra-wideband techniques have primarily targeted industrial applications due to the requirement of custom hardware at each node. Companies such as Time Domain [21] and Ubisense [81] are the major players in the UWB indoor localization space. Due to the custom hardware utilized in these systems, per-unit prices are high, limiting their applicability in consumer-facing application domains.

The advent of commodity ultra-wideband radios has shifted the tables for the realization of ultra-wideband localization systems in consumer-oriented environments. The DW1000 IC [18], while too expensive and high-power for some applications, provides some insights on the eventual hardware requirements for an all-encompassing anchor design. The DW1000 is able to provide channel impulse response measurements, provides analyses of these measurements, and has the capability to perform data communication between radios as well.

A myriad of research has brought forth different architectures and techniques for the realization of UWB localization [24, 53], however few to none of these systems has shown adequate performance at a low price point. The following section will detail the issues

faced by the current state-of-the-art.

2.4.3 Limitations of State-of-the-Art

Despite the commercial realization of ultra-wideband localization technologies, a few limitations still exist which has greatly hindered their widespread adoption. First, the systems are predominantly expensive, hindering their adoption in residential settings. Secondly, the systems are sensitive to narrowband interference, potentially causing service outages in cases where other (licensed or unlicensed) emitters of narrowband radiation are present. Third, these systems provide little flexibility in the tag design, alienating applications which have constrained size, weight, and power (SWaP) requirements. Lastly, the systems have no support for wireless synchronization primitives which will be necessary to aid in realistic deployments.

The following chapters will detail the work performed (or proposed) to address these limitations. A deep analysis of the ultra-wideband channel impulse response and the methods to obtain its approximation will guide the design decisions made.

2.4.4 Bandwidth/Time Resolution Duality

As discussed in Section 2.4.2, a direct measurement of the ultra-wideband channel impulse response is difficult due to the limited time-domain resolution of ultra-wideband transmissions. This has lead us to pursue other methods of its capture.

One observation is that of the time/frequency duality exhibited through the Fourier and inverse-Fourier transformations [36]. This states that if a signal's frequency-domain Fourier coefficients are known, the corresponding time-domain signal can be reconstructed. So-called 'bandstitching' techniques have been used in the past [72] to cover multiple adjacent ISM bands to increase time-domain resolution. Through an extension to the UWB domain, this allows for a significant reduction in the complexity of the ultra-wideband receiver.

Furthermore, bandstitching is designed to utilize well-known narrowband receiver design techniques which are simpler and less costly to implement. The bandstitching results described in the following sections are obtained using off-the-shelf software-defined radios with little to no required hardware modifications.

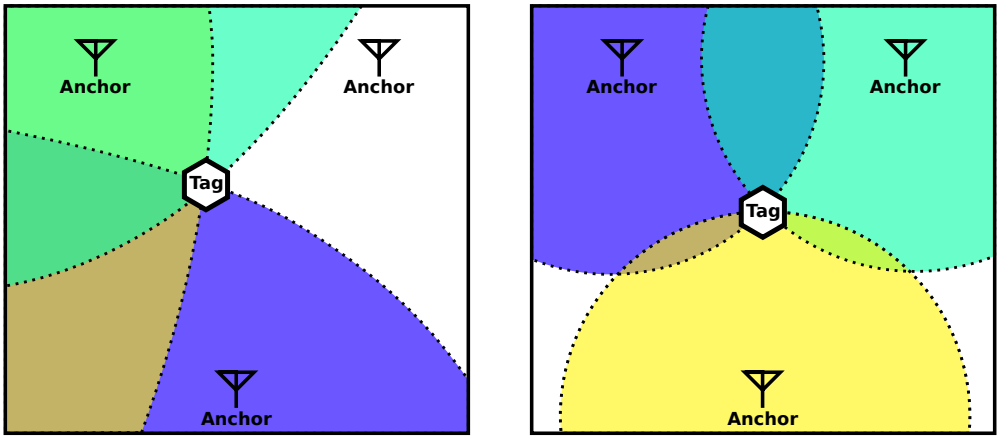

(a) TX- or RX-only Localization (Multilateration)

(b) Range-based Localization (Trilateration)

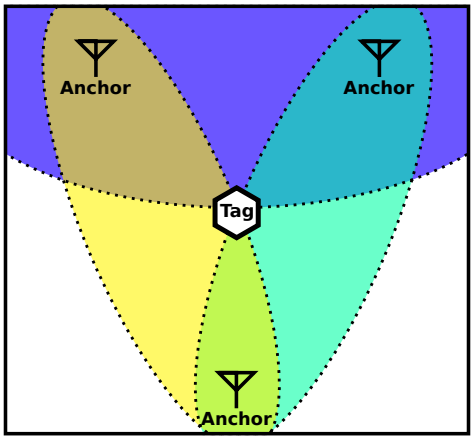

(c) Backscatter-only Localization (Ellipse-based)

Figure 2.6: **Principles of UWB Infrastructure-Based Time-of-Arrival Localization.** Infrastructure is situated throughout the observation area and used to measure the time-of-arrival of ultra-wideband signals. The timing of these signals is used to determine the position of tags using a variety of different techniques, depending on the type of tag. (a) Transmit-only tags require the use of multilateration techniques to determine the tag's location using the time difference-of-arrival of the tag's transmissions across all receiving anchors. (b) Tags which also incorporate UWB receivers are able to calculate the round-trip time-of-flight of exchanged packets with surrounding anchors to generate range estimates necessary for localization via trilateration. (c) Finally, backscatter localization schemes measure reflections of UWB signals transferred between anchors to determine a tag's location via the intersection of ellipsoids.

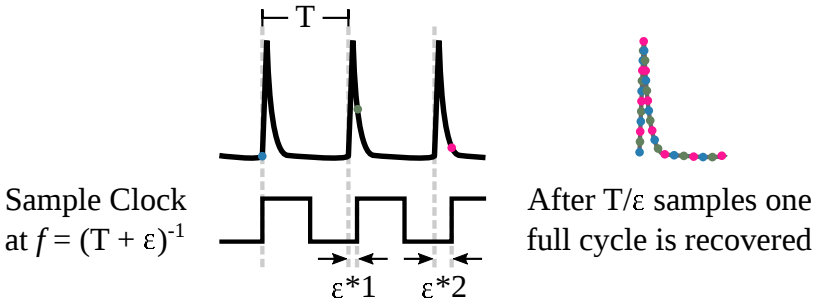

(a) Time Segmentation

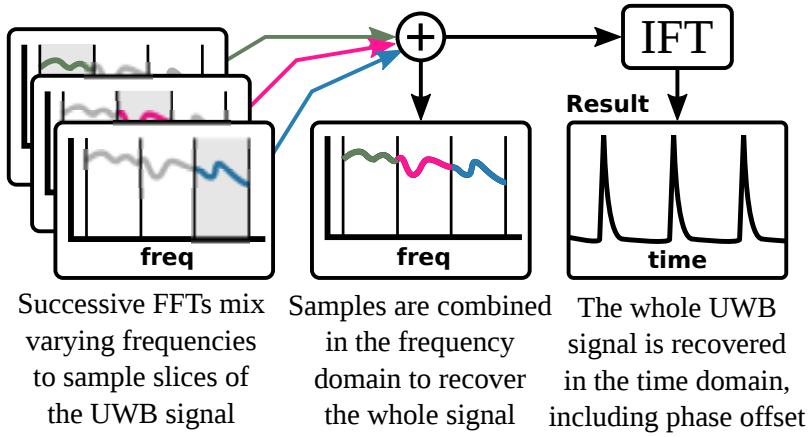

(b) Frequency Segmentation

Figure 2.7: **Reducing the Burden of UWB Reception.** Illustration of time- and frequency-segmentation techniques which can be leveraged to reduce the ADC speed required for UWB time-of-arrival estimation. By sampling at a rate just below the transmitted pulse repetition frequency, a time-stretched representation of the received signal can be reconstructed at a time resolution equivalent to a direct sampling approach. Alternatively, frequency segmentation can be used to construct the equivalent time-domain representation by successively sampling different bandwidths, stitching them together in the frequency domain, and applying the inverse Fourier transform to recover the time domain representation.

2.5 Summary

Indoor localization has exhibited a storied past, showing widely varying performance across a variety of localization media. For the purposes of discussion in this dissertation, we have chosen the RF medium, as it has shown high market penetration and high coverage in crowded, indoor environments. Within the spectrum of RF techniques, ultra-wideband methodologies have been the only techniques shown to achieve the sub-meter position determination performance required in many indoor localization applications.

Ultra-wideband technologies have not yet seen widespread adoption due to the inflexibility and tailored design of current UWB localization hardware along with the systems built around them. The following sections will explore three indoor localization applications with widely-varying tag requirements, the inadequate coverage of these applications with current solutions, and candidate solutions to the issues identified. Finally, we discuss ways in which ultra-wideband infrastructure can be more easily integrated in with existing hardware along with techniques aimed at reducing the burden of their deployment.

CHAPTER 3

SurePoint: Enabling Personal Navigation in Crowded, Indoor Environments

The realization of an accurate and responsive personal indoor navigation system will allow for the assisted navigation through indoor environments in a manner similar to that which GPS has enabled to date. Upon entering an unknown indoor environment, turn-by-turn directions to an unfamiliar destination could be given if a map of the environment and the user's location are known.

The smartphone provides a reasonable target platform for indoor use due to its pervasiveness and extensive use in similar outdoor scenarios. Smartphones use *order Watts* to power their global navigation receivers, providing a target power budget for a switch to indoor navigation. Accuracy must exceed *one meter* to achieve doorway-level resolution, and the infrastructure costs must be minimal to aid in its adoption.

The SurePoint system has been architected for the realization of accurate personal indoor navigation. In order to minimize system cost, it leverages the DecaWave DW1000 UWB transceiver, which provides communication and time-of-arrival primitives necessary for time-of-flight based localization schemes. The remaining sections in this chapter discuss the SurePoint system implementation along with design decisions which have been made to improve system localization accuracy and minimize latency.

3.1 The Flexible TriPoint Module

A standalone module was developed to provide an abstraction for localization operations. The module integrates a commodity microcontroller, the DW1000 transceiver, an RF switch, and power conditioning circuitry to provide a drop-in module for orchestrating and generating localization-centric measurements. The module is developed as a castellated package to allow for direct surface-mount soldering to a user's PCB. The following sections detail the use and necessity of each of the components used on the TriPoint module.

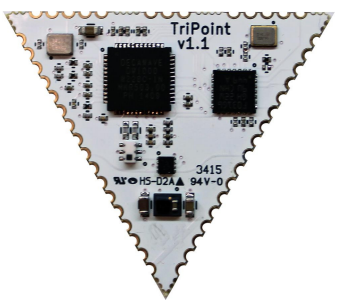

(a) TriPoint module

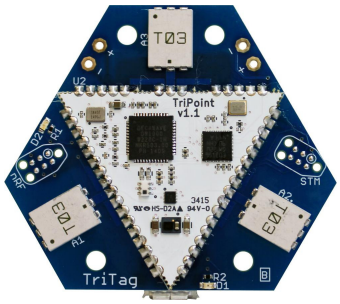

(b) TriTag carrier board

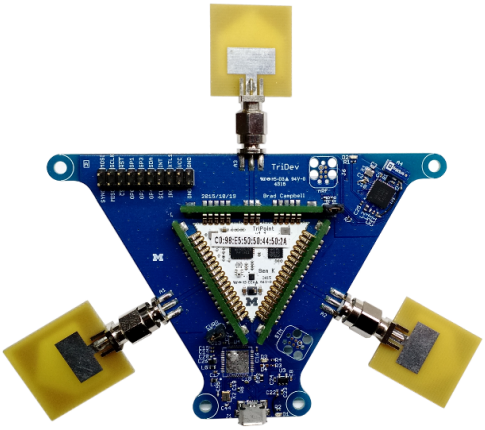

(c) TriDev Calibration Board

Figure 3.1: **SurePoint Hardware.** (a) shows the TriPoint module, a solder-on device that implements the SurePoint system and exposes location information over an I^2C interface. In (b), TriPoint is soldered onto a carrier board that includes a BLE interface. (c) shows the TriDev calibration board, allowing temporary TriTag connection to SMA antennas (pictured) or direct coaxial connection for multipath-free calibration measurements.

3.1.1 Commercial UWB Transceiver

The DecaWave DW1000 UWB transceiver provides the capability to send and receive 802.15.4a packets as well as the ability to measure the precise arrival time of these packets. During the preamble portion of each packet received, the transceiver constructs its representation of the ultra-wideband channel impulse response. From this representation, the CIR's leading edge is calculated and used as the best-guess estimate for the packet's arrival time. All packet data and timestamps are made available to an external microcontroller via a dedicated SPI bus.

Command	Description
Info	Read the TriPoint identifier and version number.
Config	Configure tag and anchor options.
Read Interrupt	Query the Harmonium for the interrupt reason.
Do Range	If not periodically ranging, initiate a ranging event.
Sleep	Put the TriPoint in sleep mode.
Resume	Wake the TriPoint and resume previous operation.
Set Location	For anchors, configure their location.
Read Calibration	Read the calibration constants for this TriPoint.

Table 3.1: **I^2C Interface to TriPoint.**

3.1.2 Dedicated MCU for Protocol Orchestration

The DW1000 UWB transceiver provides necessary but insufficient capabilities for providing position estimation. For instance, in a simple two-way time-of-flight ranging application, a total of three packets must be passed between nodes in order to ameliorate bias due to unsynchronized node clocks. For this reason, a separate microcontroller has been added to the TriPoint system to control the transceiver and orchestrate the communications required for the desired localization application.

The STM32F031G6U6 [74] was selected due to its cost (1 USD @ 3k quantity), performance (48 MHz ARM M0), and available peripherals (SPI, I2C, GPIO). The MCU orchestrates ranging operations with nearby anchors through communication with the DW1000 transceiver via SPI. A generic I2C interface is made available to provide a generic abstraction layer to the higher-level system in which the TriPoint module is placed.

3.1.3 TriPoint I^2C Interface

The TriPoint module provides an I^2C interface plus an interrupt line for configuration and location updates. A set of commands, briefly highlighted in Table 3.1, are defined as I^2C read and writes to configure and control the TriPoint module, including a command to set the node as an anchor or tag. To simplify the requirements of the application microcontroller, TriPoint is exclusively an I^2C slave. To allow TriPoint to notify the application MCU that a location is available or other event, the interface also includes an interrupt line, which when asserted the application processor can issue a "Read Interrupt" command to query the source of the interrupt.

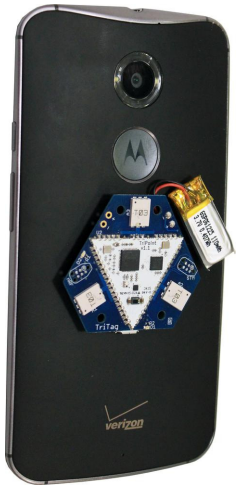

(a) TriTag with smartphone

(b) Base SurePoint smartphone UI

Figure 3.2: **TriTag Interfaces with Smartphones.** TriTag is designed to be an accessory for smartphones that provides localization for users, and the user's current position is displayed on the phone in a UI similar to the one shown in Figure 3.2b.

3.1.4 TriTag Carrier Board

TriTag, shown in Figure 3.1b, is an example of a carrier board that leverages the TriPoint module for localization. On the top side is the module plus three surface mount UWB antennas, and the reverse side of the board contains a Bluetooth Low Energy (BLE) radio. TriTag is designed to interface with smartphones to provide handheld localization for users, and Figure 3.2a shows how the system can be a smartphone accessory.

3.1.5 Antenna Diversity for Multiple CIR Observations

In order to obtain accurate time-of-flight measurements, the line-of-sight path must be observed between the nodes participating in the message exchange. Nulls in the antenna gain pattern are an inevitability, which will significantly affect ranging performance if the line-of-sight path happens to fall into one of the antenna's nulls. For this reason, the TriPoint module leverages an RF switch to select between three antennas in order to avoid the deleterious effects of antenna nulls on localization performance.

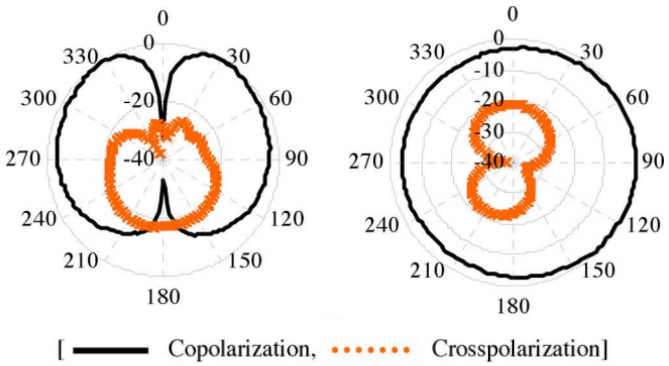

Figure 3.3: **UWB PCB Antenna Performance.** Measured antenna radiation patterns for the antenna used in our implementation at 3.4 GHz from Azim, et al. in the xz- and yz-planes, respectively [3]. Significant attenuation can be seen in cross-polarization orientations along with certain orientations in the co-polarized xz-plane, which can delay the perceived time-of-arrival for the LoS path.

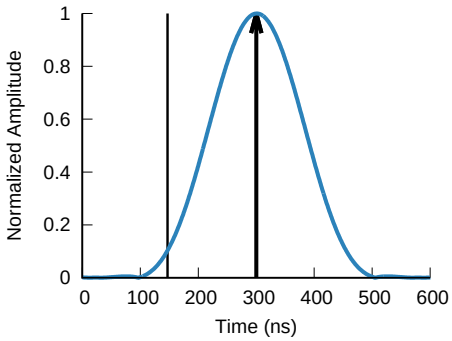

(a) One Propagation Path (Line-of-Sight)

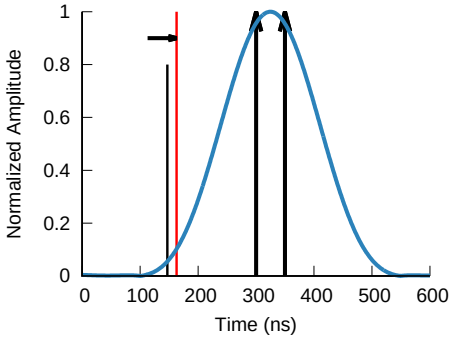

(b) Two Propagation Paths (equal phase)

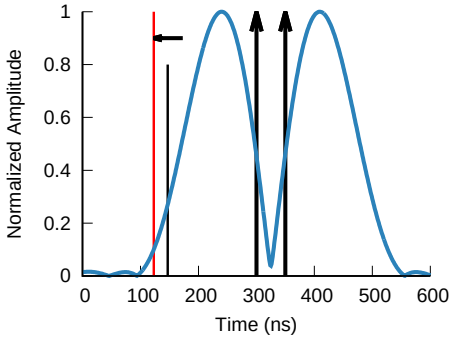

(c) Two Propagation Paths (opposite phase)

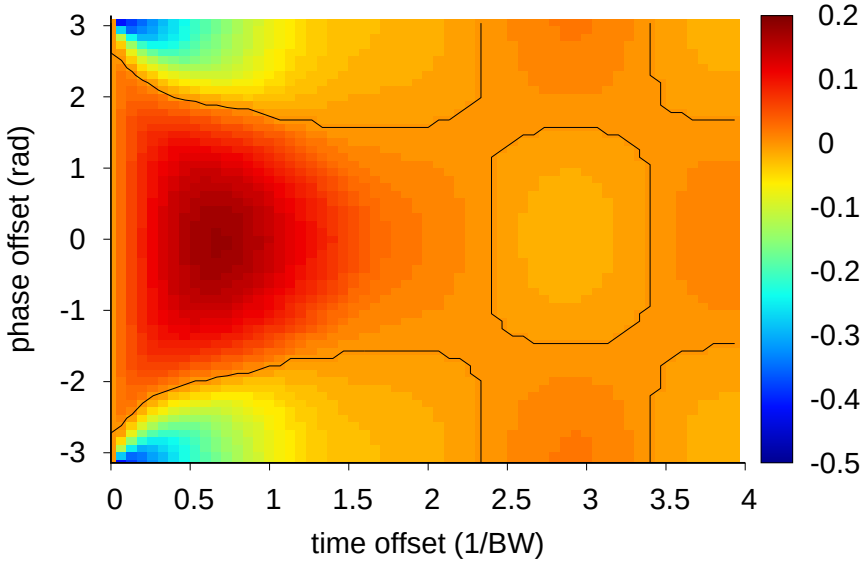

(d) ToA Error at 10% Height

Figure 3.4: **Super-Resolution ToA Distortion.** The effect on perceived time-of-arrival with two paths of equal amplitude, varying the time and phase of the second propagation path. Both positive and negative ToA bias can be observed depending on the phase and arrival time of the second propagation path. Negative ToA bias occurs from amplitude bias incurred during normalization operations applied with the chosen IFT-based super-resolution approach. Other super-resolution techniques have been proposed which achieve worse average-case ToA error [38] due to their inability to accurately fit a model to complex CIR observations.

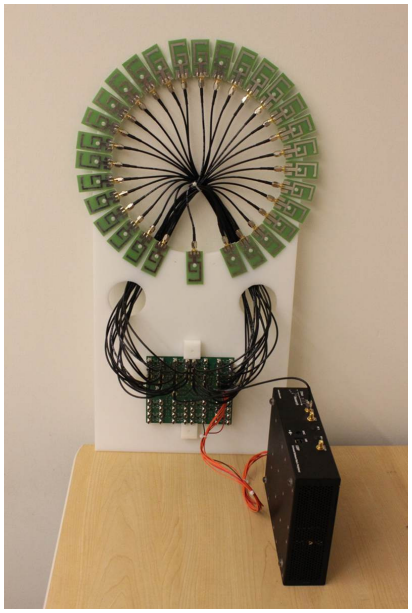

Figure 3.5: **Narrowband Diversity.** Circular switched-antenna array for spatial and polarization diversity measurements in a narrowband environment. This design is composed of a 1x64 RF switch connected to 30 dual-band PCB antennas, arranged in a circular orientation, and attached to a USRP1 platform. Beyond its use in UWB settings, antenna diversity in narrowband settings can greatly improve ToA estimation issues that plague bandwidth-limited, commodity localization system designs.

3.2 Need for Diversity in Multipath-Rich Environments

Beyond the reasons of avoiding inaccuracies due to antenna nulls cited in the previous section, additional independent observations of the RF channel can further improve time-of-arrival estimation over that of a single observation [55, 90]. Distortions in the perceived first path arrival time can be caused by multipath occurring directly after the line-of-sight path. The effects of these distortions is based on the magnitude, phase, and arrival time of the following multipath. Figure 3.4 shows the distortion in perceived time-of-arrival for a two-path scenario, varying the phase and arrival time of the second propagation path.

Additional, independent measurements can help to avoid the large deviations seen in any single range measurement. These independent measurements can come from diversity in space, polarization, or frequency. SurePoint measures range across every combination of tag antenna, anchor antenna, and RF channel available to get a total of 27 different range measurements. Assuming the range error distribution follows the approximate shape of that obtained from simulation, SurePoint takes the 10th percentile estimate over these 27 different observations to come up with the true estimate of time-of-arrival.

3.3 Efficiently Measuring Many, Independent Range Observations

A naïve approach to obtaining these 27 different observations would require the exchange of $27 \times 3 = 81$ packets assuming a typical two-way time-of-flight exchange for each observation. This has a significant impact on the speed and latency of the system's location estimation capabilities, so a custom protocol has been designed to improve the maximum attainable update rate without sacrificing localization performance.

Figure 3.6 shows the protocol employed by SurePoint to minimize the number of packet transmissions, and therefore maximize the system's maximum attainable position update rate. The protocol starts with a series of 27 one-way transmissions across each independent antenna and channel combination.

Next, a two-way time-of-flight exchange is performed to calculate the clock offset between the tag and any nodes participating in ranging operations. This consists of a final one-way transmission followed by a contention window for responding nodes. The relative time-of-arrival offset between each of the preceding one-way ToA measurements can then be back-calculated based on the newly-synchronized timing information.

3.3.1 DW1000 Calibration

One of the most challenging aspects of the DW1000 is that each chip requires calibration of both the intrinsic transmit delay and receive delay that is custom to each chip [17]. Unfortunately, some settings such as the transmit power level change this delay constant, so calibration must be collected for all possible radio configurations.

In order to produce an accurate calibration estimate, any deleterious effects of multipath must be sidestepped to compare the time-of-flight delays to a known, distortion-free range. This either requires the use of an anechoic chamber (*costly*) or a production-time hard-wired calibration procedure. Harmonium utilizes a production-time calibration procedure to measure the RX and TX delays of each TriPoint module through temporary use of the TriDev board, which allows for direct-to-SMA characterization measurements.

To facilitate calibration, we design the TriDev board with spring-loaded pins. The castellated headers allow simply inserting the TriPoint module, running calibration, and recording it to the onboard flash memory. A central database of calibration values is kept for the purposes of reprogramming at a later date.

The calibration procedure starts with the production of a reference node with known TX+RX delays for each channel. This is produced in a round-robin manner with two other TriPoint modules. Range estimates are produced for each pairing of the three nodes (A to B, B to C, A to C). The true range (including all cabling) is subtracted from these range estimates. This gives a series of linear equations ($t_{AB} = t_{A,error} + t_{B,error}$, $t_{BC} = t_{B,error} + t_{C,error}$, $t_{AC} = t_{A,error} + t_{C,error}$) which can be solved for $t_{A,error} = (t_{AB} + t_{AC} - t_{BC})/2$. $t_{A,error}$ consists of the sum of the transmit and receive delays for each of the reference node's channels. Subsequent nodes can compute $t_{X,error}$ from any additional error observed while performing calibration measurements with the reference node.

These measurements are sufficient for single-channel two-way time-of-flight calibration. However, for the Harmonium ranging protocol, these time delays must be broken out into RX and TX delays for each channel independently. This would require the measurement of true TX and RX delays for the reference node. However, this is not possible without the use of a high-speed oscilloscope and fine-grained introspection into the DW1000's inner architecture. Instead, we choose to again reference each channel's RX and TX delays to the reference node by setting the reference node's transmit delays to zero. This results

in N^2 pairings of RX+TX delays for each {TX channel, RX channel} combination

$$\begin{bmatrix} t_{A1}+r_{A1} & t_{A2}+r_{A1} & t_{A3}+r_{A1} \\ t_{A1}+r_{A2} & t_{A2}+r_{A2} & t_{A3}+r_{A2} \\ t_{A1}+r_{A3} & t_{A2}+r_{A3} & t_{A3}+r_{A3} \end{bmatrix}$$

which is used as the true calibration data loaded onto each node. The arbitrary TX delay = 0 setting for the reference node introduces

$$\begin{bmatrix} 0 & t_{R2}-t_{R1} & t_{R3}-t_{R1} \\ t_{R1}-t_{R2} & 0 & t_{R3}-t_{R2} \\ t_{R1}-t_{R3} & t_{R2}-t_{R3} & 0 \end{bmatrix}$$

of error to each node's calibration which gets factored out when performing two-way ranging between two nodes calibrated to the same reference node, as the full two-way ranging calculation uses a sum of elements across the diagonal. These calibration values are stored in the TriPoint module's flash, hiding this implementation complexity. This system is robust to loss or failure of the reference node as well, as a second reference node can be recreated from any two nodes which were previously calibrated to the original reference.

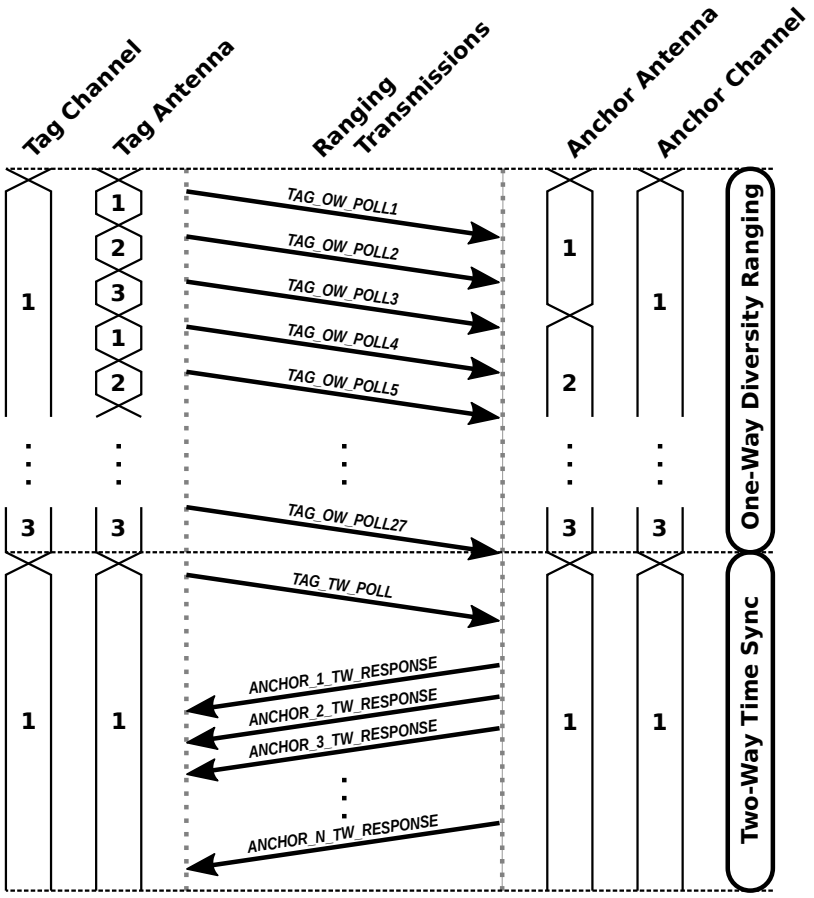

Figure 3.6: **SurePoint Ranging Protocol.** SurePoint is able to leverage antenna diversity without significantly impacting position update rate. The protocol starts with a series of 27 broadcast transmissions from the tag for each combination of tag antenna, anchor antenna, and RF channel. The inter-packet time-of-arrival data collected from this sequence provides information on the difference between all range estimates throughout the sequence. Finally, a two-way time-of-flight handshake is sent to determine the true range estimate for the first configuration—a total of *28* POLL messages. The offset between the first and last poll message is used to calculate the crystal frequency offset between the tag and anchor. From the first measurement, the time offset between the tag and anchor is known, leading to estimates of range for the other 26 combinations from the initial difference-based measurements.

3.4 From Ranging to Localization

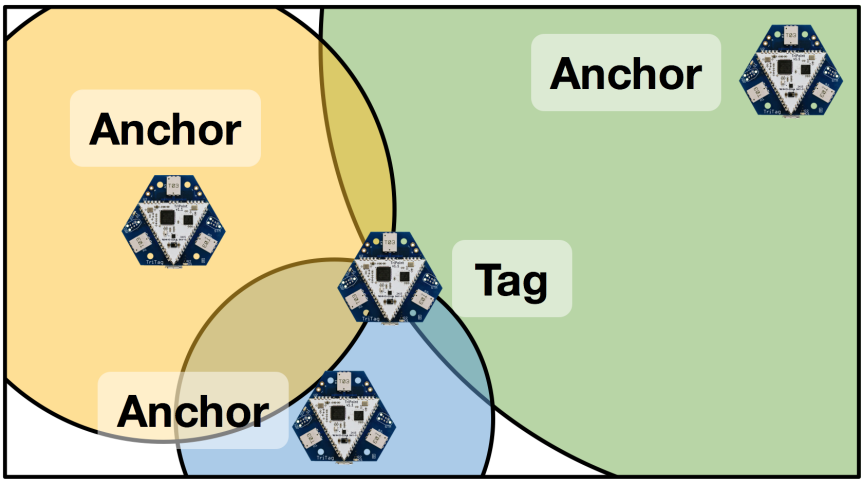

Figure 3.7: **2D Trilateration.** In two dimensions, each range estimate between tag and anchor gives a circle of possible locations. With two anchors, the circles intersect in two locations, requiring a third anchor to disambiguate the tag position. In general, trilateration requires one more anchor than the number of dimensions. Thus to resolve position in 3D, SurePoint requires range estimates to four anchors.

To localize nodes, SurePoint measures the pairwise distance between the mobile tags and a network of anchors with known, fixed locations. From our previous discussion on the impact of diversity, however, we find that range estimates improve significantly with multiple, diverse measurements. Doing a full diversity handshake between the tag and each anchor is prohibitively expensive. Instead, in this section we describe an efficient broadcast ranging protocol that adds minimal overhead for every additional pairwise range.

3.4.1 SurePoint Ranging Protocol (Figure 3.6)

Per the discussion in section 3.2, Harmonium requires the calculation of 27 range estimates *per anchor*. If one were to calculate range estimates using a traditional two-way time-of-flight (TW-ToF) approach, this would result in a minimum of $27 \times 4 \times 2 = 216$ *packet transmissions*.

In order to reduce the number of packets required to leverage the additional diversity, SurePoint developed the original broadcast ranging protocol which we describe here [44]. We improve upon this protocol by adding much-needed robustness. In our original SurePoint experiments, a tag is placed in an environment with 16 anchors, yet only hears from an average of 4 each round and 9 in the best round. In contrast, SurePoint hears responses from 7.3 of the 9 deployed anchors on average.

Harmonium applies two optimizations to greatly reduce the number of transmissions required to integrate antenna and frequency diversity between Harmonium nodes. The first optimization aims to reduce the number of required tag to infrastructure transmissions. A naïve implementation would require 27 $\times$ N unicast transmissions between tag and infrastructure to satisfy the poll phase of the traditional two-way time-of-flight ranging protocol. This is instead grouped into a single sequence of 27 *broadcast* transmissions from the tag which are unequivocally received by all nearby infrastructure.

The second optimization utilizes the properties of reciprocity to reduce the number of responses from each anchor from 27 to *1*. The time difference of the remaining 26 is already known from the initial observed broadcast sequence. The received time-of-arrival is instead contained in this final ranging response packet for later back-calculation. The antenna which is used by each anchor in sending their response is chosen heuristically to be the antenna which received the most number of packet broadcasts from the tag during the ranging broadcast sequence round.

At this point in the original SurePoint protocol, a single packet duplicating the original packet is sent in fixed time slots from each anchor. This duplicate packet is critical as it is used to calculate the crystal offset between the tag and the anchor. The original design has two major disadvantages. The first is that every anchor is assigned one of sixteen (the number of deployed anchors), fixed global time slots at compile time. This inflexibility is intractable for real-world deployments. Secondly, as crystal offset compensation is critical to capturing a range estimate, missing either the first or last packet would invalidate the entire ranging sequence.

The SurePoint protocol replaces the static anchor assignments with three contention windows to provide anchor responses to the tag. The anchor chooses a random time slot within each to respond to the tag and continues to transmit its responses until an acknowledgment is received back from the tag signaling a complete packet reception.

Furthermore, we add two additional crystal calibration packets, now sending one on each channel. From a modest empirical study, we find that leveraging frequency diversity over antenna diversity provides the greatest likelihood that at least one calibration packet is received successfully.

While these revisions resolve the limitations imposed by the static assignment and selection of anchors, they do not solve the principle scalability challenge, coordinating multiple tags such that they do not interfere with one another. As the length of the ranging sequence is quite large, coordinated timeslots are required to efficiently support a large number of tags with a high update rate. We address multi-tag coordination in Section 6.4.

3.4.2 Position Solver

In order to determine the tag's location, Harmonium leverages trilateration, which calculates an object's position using its range to a number of known-location objects. These range measurements form a number of equi-distant surfaces around each known-location object which describes the set of possible locations which are consistent with all measured range estimates. In two dimensions, the intersection of two circles gives up to two points, which requires one more range estimate to disambiguate the object's correct two-dimensional location. Still one more range measurement (four total) is required to trilaterate an object's position in three dimensions.

Although four range estimates are required to uniquely determine an object's position in three dimensions, even more range estimates help to improve the reliability of the resulting position estimate. If more than four ranges are given, a unique solution may not be possible given any error in range estimates provided. SurePoint adopts the insights from Gezici et al. [29] and adds iterative point removal to a non-linear least squares solver [92] to trilaterate. The iteration enables SurePoint to eliminate ranges which are likely causing the greatest amount of error.

Once a SurePoint tag has all of the range estimates, SurePoint solves for position using a simple minimum mean squared error algorithm. SurePoint optimizes using the quasi-Newton BFGS method as it performs well and converges relatively quickly [92]. Guided by Gezici's observations [29], we enhance the original SurePoint solver by adding an iterative step that considers removing an anchor whenever the error fails to converge to better than 10 cm.

3.5 SurePoint Performance Assessment

We begin our evaluation of SurePoint with two experiments evaluating the quality of SurePoint as a localization system. We then use these traces to explore the protocol decisions made by SurePoint. Next we examine the performance of SurePoint in multi-tag settings. Finally we consider some microbenchmarks for the TriTag modules.

3.5.1 Figure 3.8: Stationary Tracking

We begin by replicating the cross experiment from SurePoint in a near-identical evaluation space.[1] In this experiment, we stand for approximately 15 s at 50 positions in a cross in the

[1] A 20×20 m room. However our room also includes furniture, persons playing ping-pong, and other perturbations not present in SurePoint's evaluation space.

middle of the evaluation area. For our experiment, we deploy on 9 anchors as opposed to SurePoint's 15 anchors. Despite this, our system hears from an average of 7.3 anchors over the course of the entire experiment as opposed to SurePoint's 4, demonstrating the efficacy of our reliability improvements.

As a consequence of the improved reliability, we show a significant improvement in positioning accuracy, especially in the long tail where where are able to improve the 99th percentile accuracy from 2.02 m to 0.77 m. While we generally eschew interpolation of points, leaving that for higher-level systems and aiming to provide better raw data to them, we process the location estimates from this experiment with a simple 3-point median filter (i.e. reject a single outlier), in which case the maximum error is only 0.76 m.

The tag will only attempt to range with the anchors if it hears the scheduling packet. Over the course this experiment, the tag was scheduled 840 times and the tag successfully participated in in 837 rounds (99.6%). Within each round, this was the only tag in the environment, and the tag was thus eligible to range during all 12 slots. Of the 10,044 slots that the tag was aware it was eligible to range, it successfully recovered a range estimate in all 10,044.

3.5.2 Figure 3.9: Tracking Motion

The number of packets required to capture full diversity data extends the duration of a single ranging event to about 80 ms. One natural question then is whether this makes SurePoint susceptible to motion blur or other errors introduced by tracking fast-moving objects. In Figure 3.9, we find that SurePoint can track an object moving at up to 2.4 m/s without degradation in tracking quality.

As a trilateration-based system, SurePoint benefits from diverse anchor placement. However, the nature of building construction is such that it is much easier to achieve large diversity in X and Y, but less so in Z. We calculate the tracking accuracy for each major component and find that the median X error is 0.06 m, the median Y error is 0.07 m, while the median Z error is 0.15 m, nearly triple.

3.5.3 Diversity

SurePoint goes to great lengths to capture large amounts of diversity to improve the range error. As the ultimate goal is location, one reasonable question may be to ask whether it was all worth it. Figure 3.11 explores what happens if one eliminates some diversity by reprocessing a data trace and dropping measurements. A roughly 9 cm median increase in range error is expected, however, these ranges are then fed into a solver, which may be able

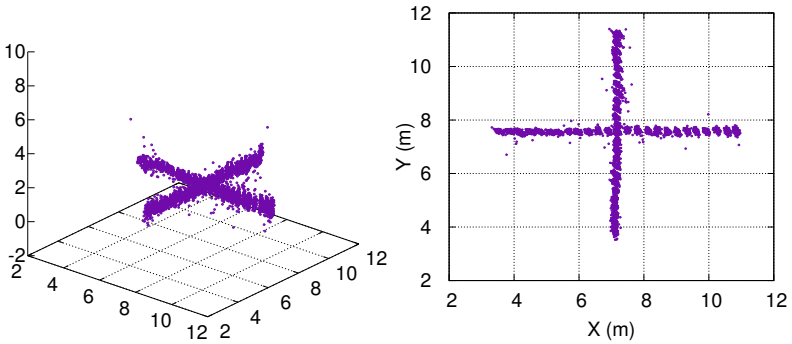

(a) Profile View

(b) Top View

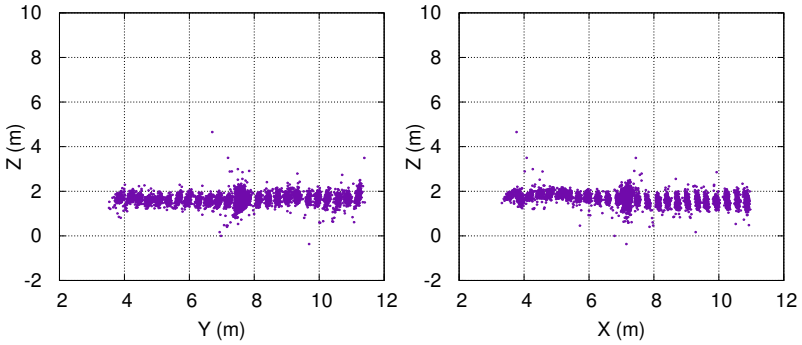

(c) View from X

(d) View from Y

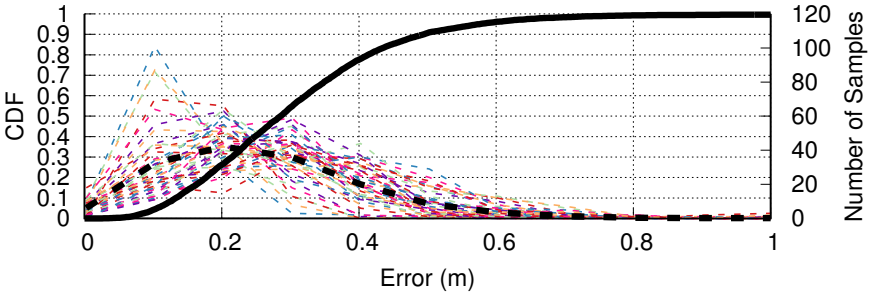

(e) CDF of Aggregate Error and Histograms of Point Errors

	Min	**50%-ile**	**90%-ile**	**95%-ile**	**99%-ile**	**Max**
Accuracy	0.01	0.29	0.50	0.58	0.76	2.35
Precision	0.01	0.12	0.28	0.35	0.55	0.93

(f) Figures of Merit

Figure 3.8: **SurePoint Stationary Tracking Experiment.** To evaluate the baseline localization accuracy of SurePoint we place a node at 50 locations for 15 s each. Each point is 1 foot apart (the spacing of floor tiles in the evaluation area). Across this entire sample set, SurePoint achieves 0.29 m median accuracy, 0.12 m median precision, and is further able to realize sub-meter 99th percentile accuracy and precision.

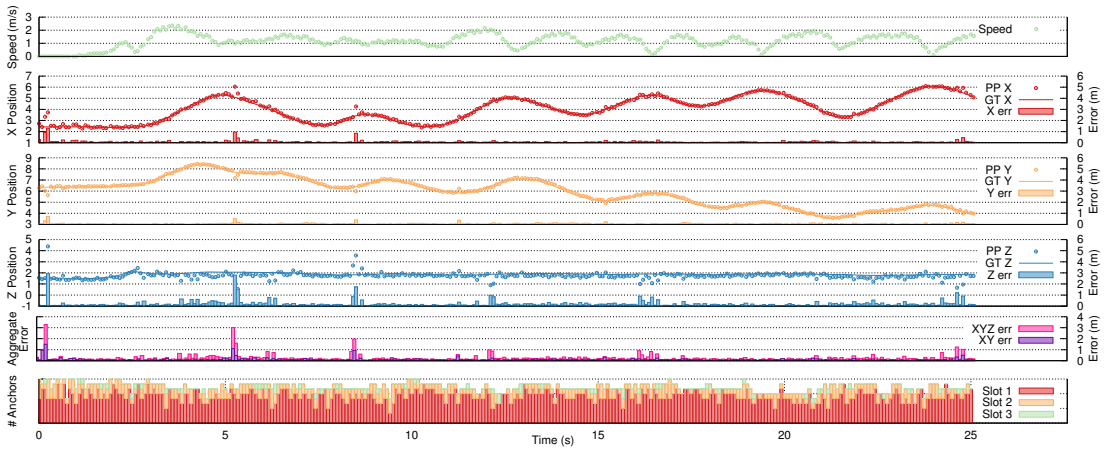

Figure 3.9: **SurePoint Mobile Tag.** The addition of diversity to the ranging protocol means that capturing a single sample spans about 30 ms. In this experiment we move a tag throughout the environment to evaluate whether this motion has a negative impact on SurePoint location estimates. We find no correlation between the speed of the tag and the instantaneous error, even at 2.4 m/s. We explore this experiment further by separating out the tracking quality in the X, Y, and Z dimensions. The nature of physical buildings allow for much greater diversity for anchor placement in X and Y compared to Z. As a result, Z error contributes the most to overall error. Finally, to give an intuition for the impact of the anchor reply contention, we show the distribution of anchors that replied in each slot over time.

to overcome some of the additional noise in the input data. In practice, we find that the solver alone is not sufficient to overcome the extra error introduced by the noisier range estimates, and the full diversity performs best.

Another interesting question to ask may be whether we could reduce the diversity overhead for short periods of time, exploiting some spatial or temporal case where certain channel/antenna configurations consistently return the best range estimates. Figure 3.10 digs into the ranging performance at two nearly identical locations in space, sampled one after another in time, and finds that there is not a consistently good (or bad) set of configurations that SurePoint could temporarily restrict itself to measuring

Indeed, this result is not surprising when one considers the sensitivity of the UWB channel impulse response. During other testing, we inadvertently replicated the results from Adib et al. [2], detecting the respiratory rate of the experimenter. As we collected this data in a busy environment, there is little surprise that no exploitable spatio-temporal efficiencies could be found.

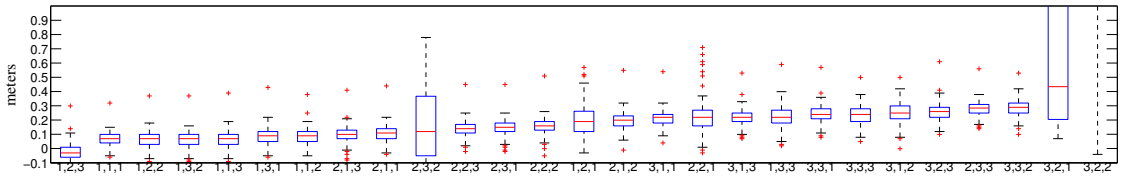

(a) Raw range distribution error from Figure 3.8 cross center to Anchor 2B, sorted by median error.

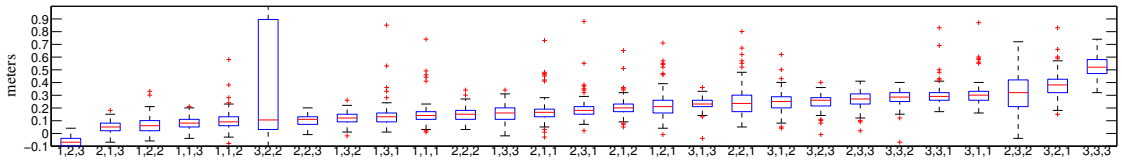

(b) Raw range distribution error from 0.3 m (1 ft) from Figure 3.8 cross center to Anchor 2B, sorted by median error.

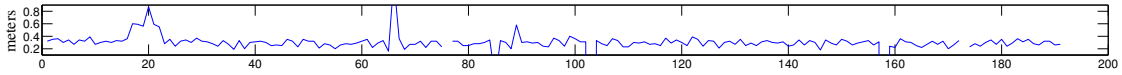

(c) Range error over time for configuration `Channel 1, Anchor Antenna 2, Tag Antenna 2` from (b).

Figure 3.10: **Does SurePoint Need all the Diversity?** SurePoint captures 27 range estimates every ranging event. Here we dig into the ranging performance for each combination of `(channel, anchor_antenna, tag_antenna)`, for two nearly adjacent points from the stationary cross experiment (Figure 3.8) to see whether there is an exploitable spatio-temporal correlation of the best (or worst) configurations. Between (a) and (b), the tag moves only 0.3 m, however the distribution of errors changes significantly. Drilling further into one configuration, (c) shows that the instantaneous error consistently varies around 0.1 m or more between measurements. From this we conclude that there is no reasonable means to reduce the number of ranges taken each round without sacrificing accuracy.

3.5.4 Anchor Response Quality

The protocol improvements introduced by this work improve upon the average number of responding anchors from Kempke's original, even given a smaller number of deployed anchors. While four replies are theoretically sufficient to precisely localize a tag, additional range estimates are preferred to refine the position estimate. Figure 3.12 explores what happens to the motion trace as anchors are artificially removed. To minimize the bias from anchor selection, we randomly select a new set of anchors to remove at each sampled point. Median error improves quickly with the first one or two additional anchors, but then tapers off. The addition of more anchors, however, continues to improve the performance of the long tail of range estimates. The experiment did not capture enough samples with 9 anchors responding to be meaningful.

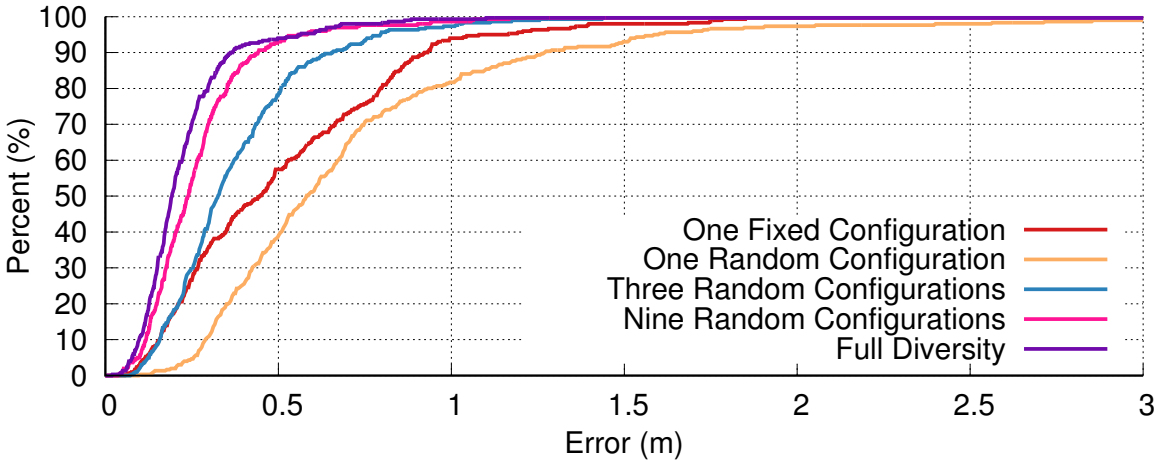

(a) CDF of Range Errors

Limit	Min	50%-ile	90%-ile	95%-ile	99%-ile	Max
One Fixed Conf.	0.06	0.45	0.92	1.08	1.76	1.88
One Random Conf.	0.10	0.58	1.27	1.59	2.86	4.21
Three Random Conf.	0.07	0.33	0.65	0.80	1.18	2.47
Nine Random Conf.	0.04	0.24	0.44	0.56	1.06	1.26
Full Diversity	0.02	0.19	0.36	0.56	0.86	1.55

(b) Figures of Merit

Figure 3.11: **Impact of Diversity on Location.** To investigate the importance of diversity for localization quality, we take the motion trace from Figure 3.9 and re-run localization processing with fewer ranges. We consider the ranges from one fixed configuration (`Ch 1, AncAnt 1, TagAnc 1`) as well as selecting 1, 3, or 9 ranges at random from the full set of diversity measurements. As the amount of diversity increases, so too does the location accuracy.

3.5.5 Iterative Solver

As our final improvement over SurePoint, we add an iterative solver that attempts to remove an anchor whenever the position fails to resolve to less than 10 cm of error. As iteration is relatively computationally expensive, in Figure 3.13 we examine the benefit that the iterative solver provides. While we see only a 0.01 m improvement in median error, iteration does particularly well in improving the long tail, reducing 99th percentile error from 1.63 m to 0.86 m. This result aligns well with expectations. The iterative solver sees the highest utility when the position estimate is poor, and for some cases it is able to identify and remove a bad range.

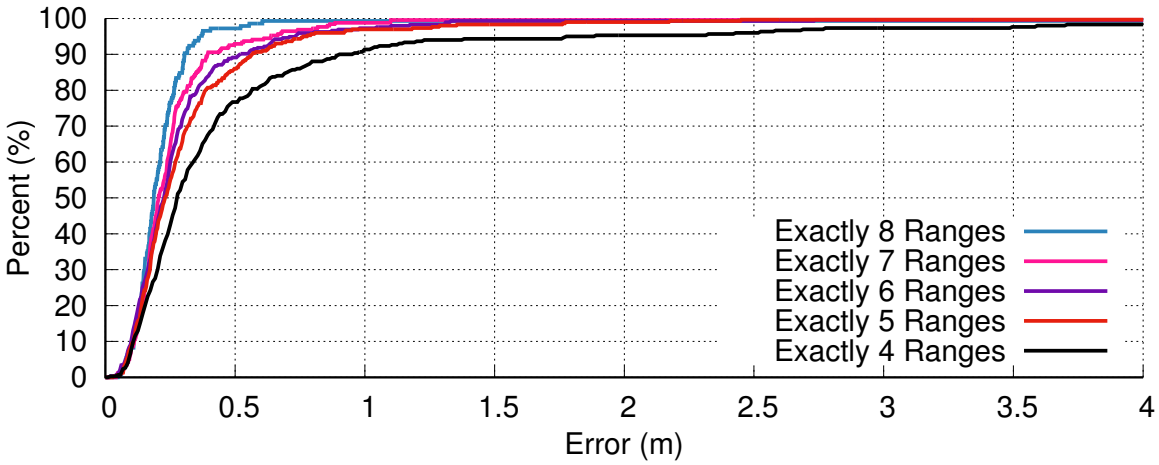

(a) CDF of Range Errors

Limit	Samples	Min	50%-ile	90%-ile	95%-ile	99%-ile	Max
4 Ranges	301	0.01	0.28	0.89	1.77	4.99	9.60
5 Ranges	301	0.03	0.23	0.56	0.76	1.77	3.57
6 Ranges	294	0.02	0.23	0.52	0.71	1.31	5.10
7 Ranges	255	0.04	0.21	0.39	0.63	0.98	1.55
8 Ranges	138	0.05	0.19	0.31	0.36	0.58	0.62

(b) Figures of Merit

Figure 3.12: **Value of Multiple Anchors.** The broadcast ranging protocol is designed to make ranging with many anchors efficient, however, trilateration only requires four ranges to resolve position. To measure the value of anchors beyond the minimum, we take a single dataset and randomly remove ranges. After one or two additional anchors, the median error plateaus, however additional ranges continue to improve the long tail performance. In some cases, more ranges may mean more outliers, accounting for the inconsistent trend of the worst-case sample as more ranges are added.

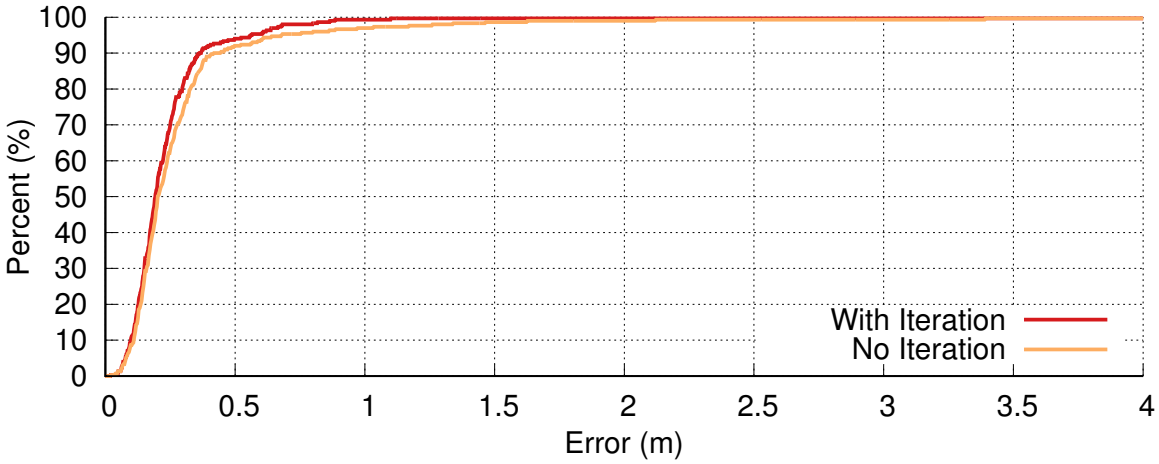

(a) CDF of Range Errors

Limit	Samples	Min	50%-ile	90%-ile	95%-ile	99%-ile	Max
Iteration	301	0.02	0.19	0.36	0.56	0.86	1.55
No Iteration	301	0.03	0.20	0.42	0.68	1.63	3.48

(b) Figures of Merit

Figure 3.13: **SurePoint Position Solving.** Guided by Gezici [29], SurePoint implements an iterative position estimator that considers dropping a range if doing so significantly (thresholded to 10 cm) improves the error from the least squares solver. This is somewhat computationally expensive, however, so here we compare against the baseline of simply including all ranges. From this dataset, iteration dropped 202 of the 4788 ranges (4%), which improved 99th percentile error by 0.77 m, a nearly 50% improvement in long-tail error, but only improved median error by 0.01 m.

3.5.6 Multiple Tags

As the precise time scheduling provided by our LWB-variant isolates operations for each ranging tag, the absolute error of each sample taken by each tag is unaffected as the number of tags increases. As the number of tags increases, however, the available sampling rate decreases, which can hinder the tracking of fast-moving objects. Given values for the round period, T_{round}, the contention and scheduling windows t_{flood}, and the duration of a ranging event, t_{slot}, the number of available slots n_{slots} can be expressed as: $n_{slots} = \left\lfloor \frac{T_{round} - 2 \cdot t_{flood}}{t_{slot}} \right\rfloor$ Our implementation currently assigns $T_{round} = 1\,\text{s}$, $t_{flood} = 10\,\text{ms}$, and $t_{slot} = 80\,\text{ms}$, resulting in $n_{slots} = 12$. While a longer round period does make available more slots, it does so at the cost of latency for new tags joining the network as scheduling events occur less often.

The other interesting side-effect of the scheduling protocol is that the update rate of even a single tag is not consistent, as the ranging tag(s) must periodically wait for a scheduling and contention slot to pass. In practice, this means our single tag-case has an effective average update rate of $(\frac{11 \times 80\,\text{ms} + 1 \times 120\,\text{ms}}{12})^{-1} = 12\,\text{Hz}$. For two tags the update rate is 6 Hz, three tags is 4 Hz, four tags is 3 Hz, and at five tags the first two tags will achieve a 2.4 Hz update rate, while the next three tags will only realize a 2 Hz update rate. This configuration allows for a maximum of 12 tags in one space, however the master could easily extend T_{round} on demand, at the expense of join latency.

Another possible point of contention is when tags join. Unlike the original LWB, which reduces the scheduling interval after an initial startup phase to save energy, SurePoint capitalizes on its wall-powered anchors to always maintain high-frequency scheduling, facilitating fast joining at any time. Figure 3.14 examines the impact of the shared contention slot when 10 tags are all powered on at the same time. Across four trials, only 3 contention slots exhibit true congestion, scheduling no tag when tags were attempting to join.

3.5.7 Microbenchmarks / TriPoint Evaluation

As a modular design, the TriPoint module adds some burden to the systems in which it is deployed. First, the TriPoint module is 5 cm^2 in size, requiring at least this much PCB area on the carrier board to which it is soldered.

In dollars and cents, the module adds the marginal cost of an additional PCB fabrication, around $0.60 at modest volume pricing. However, this has the possibility of resulting in a net cost savings at the carrier board would no longer require the use of a higher-technology four layer PCB design. The dedicated STM32 microcontroller introduces an additional cost of 1 USD in modest volume.

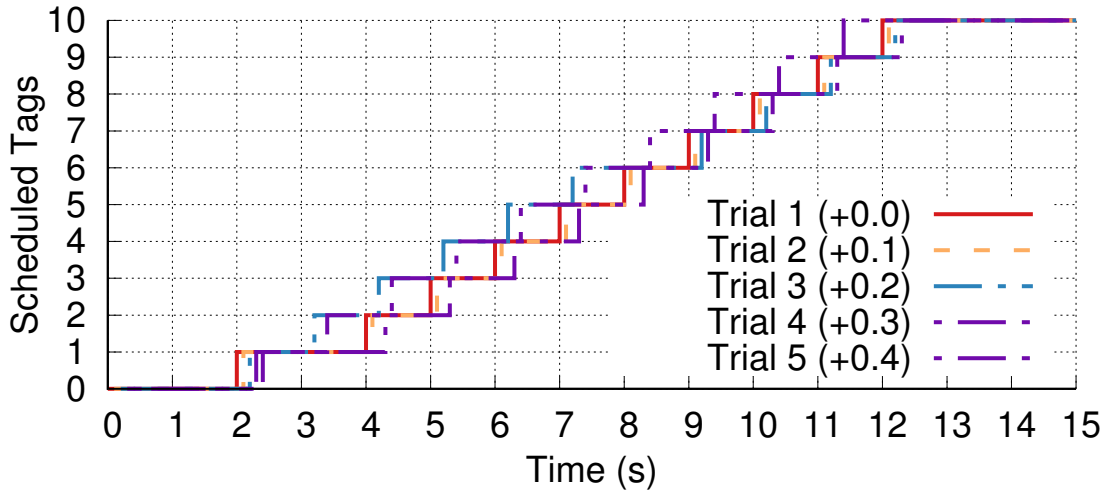

Figure 3.14: **SurePoint Tag Join Contention.** To evaluate how our system handles multiple tags entering the space concurrently, we attach 10 tags to a power strip and power all of the tags on at once. For clarity of viewing, we offset the trials by 0.1 s. In Trial 5, all 10 tags joined in 10 s, the fastest possible. Trials 1, 2, 3, and 4 all had one window where interfering tags collided in a manner such that no tag successfully scheduled.

Lastly, the power burden added by the dedicated STM32 microcontroller is approximately 18.9 mA while active. This is, however, much smaller than the DW1000's 145 mA active current while in receive mode. When the DW1000 is not active, the STM32 can be powered down, in which case it draws around 2.3 µA.

3.6 Summary

In this chapter, we have introduced the SurePoint localization system, a platform which allows for the realization of personal tracking and navigation through indoor environments. Personal navigation systems require a localization accuracy of around 1 meter in 2D environments, and necessitate an update of around 1 Hz in order to provide timely feedback to end-users of their position. With an intended integration into mobile electronics such as smartphones, an active power of one Watt, similar to GPS, is targeted. SurePoint has been designed to meet these requirements through leveraging commercial UWB transceivers. A number of enhancements were required to meet these specifications, yielding the TriPoint module which combines the UWB transceiver, an RF switch for antenna diversity, and a dedicated microcontroller for protocol orchestration.

SurePoint utilizes the same module design for infrastructure which is placed throughout the indoor environment with its location determined a-priori. Tags query range from

the surrounding fixed infrastructure and determine their position based on range estimates across a variety of antenna combinations to ameliorate the effects of multipath propagation. This tag architecture achieves the desired performance requirements in indoor settings, but does so at the cost of high tag power and complexity. In the following chapter, we will explore the application of guiding quadrotors in indoor environments, an application which necessitates a careful consideration of tag size, weight, power, and cost.

CHAPTER 4

Harmonium: Exploring Hybrid Techniques to Address System Deployability

The SurePoint system described in the previous chapter has been shown to provide adequate localization performance to aid in the realization of personal navigation and certain robotic navigation applications, yet there are a number of characteristics of the system which make it undesirable in certain conditions and applications. Over the next two chapters, we identify these pain points and propose enhancements which aim to pare down the limitations addressed.

4.1 Limitations of SurePoint Deployability

In evaluating the deployability of the SurePoint system, two problems arise. First, the system can *easily fail to operate in the presence of interfering transmissions*. Second, the tag hardware is *large*, *high-power*, and *costly*, limiting its usefulness in the asset tracking and aerial robotic tracking domains. The following sections will provide solutions to circumvent these issues in the future.

4.2 Bandstitching Enables Narrowband Rejection

Prior work [72] has shown that the effective time resolution attainable by narrowband receivers can be substantially increased through he use of a technique known as "bandstitching." Bandstitching sweeps the narrowband radio across a contiguous set of bandwidths in order to obtain an equivalent time-domain representation of a wideband signal. The following sections will detail the benefits of using bandstitching approaches in realistic indoor RF settings along with the design of hardware necessary to make these measurements.

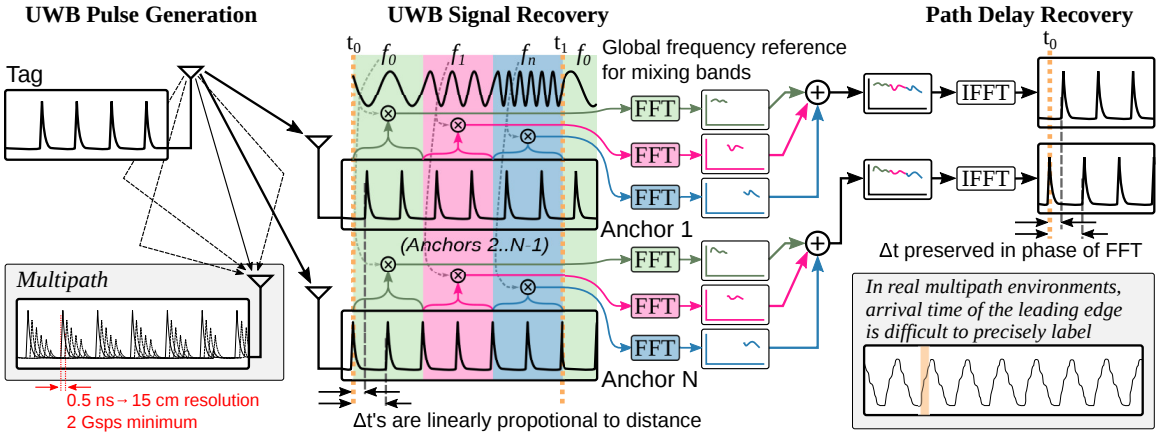

1. **Pulse Generation.** The tag UWB pulse generation design is covered in section 4.3.1.

2. **UWB Signal Recovery.** Harmonium anchors employ a novel UWB bandstitching variant, presented in section 4.2.2, to recover the transmitted signal.

3. **Path Delay Recovery.** Post-processing exploits the frequency diversity enabled by the UWB signal to extract time-of-arrival offset of the pulse train at each anchor node as presented in section 4.2.1.

Figure 4.1: **Harmonium Overview.** A mobile tag in free space broadcasts a UWB signal that is captured by anchor nodes. To localize the tag, at least four anchors must capture the tag's signal and determine the relative delay from the tag to each anchor. In complex indoor environments, reflections due to multipath make precisely identifying the arrival time difficult. To achieve 15 cm resolution, direct time-domain UWB recovery would have to sample at 2 Gsps or faster. In contrast, Harmonium adapts bandstitching to recover UWB signals, using frequency-stepped commodity narrowband RF frontends to capture successive chunks of the UWB frequency components. These chunks are combined in the frequency domain to recover the whole signal, and returned to the time domain to find the arrival time at each anchor. This approach encodes the time domain difference in arrival times at different anchors in the phase of the complex coefficients of the Discrete Fourier transform; if a signal is delayed by D samples at one anchor with respect to another anchor, then each complex coefficient of the FFT is multiplied by $e^{\frac{-j2\pi kD}{N}}$, where k is the FFT coefficient index and N is the size.

4.2.1 Measuring Time with UWB Pulse Trains

Recall that the goal is to precisely capture the time between when a signal is sent from a transmitter and is received by an anchor. There is no synchronization between tags and anchors, which means that tags cannot simply send a single pulse. Rather, tags send a continuous pulse train and anchors compute a phase offset from a shared, global time reference. Since the signal is periodic, this offset will alias if the path delay is longer than the period. In practice, this means that the pulse repetition frequency defines the maximum distance between a tag and anchor that Harmonium can measure. While the interval

between pulses can be extended to improve range, this reduces receiver SNR, affecting system performance.

4.2.2 UWB Receivers are Susceptible to Narrowband Interference

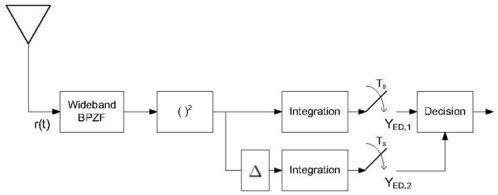

(a) UWB RX w/ Energy Detection [67]

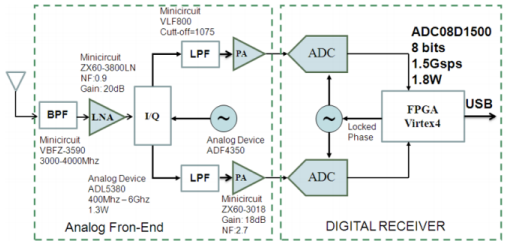

(b) UWB RX w/ ADC [73]

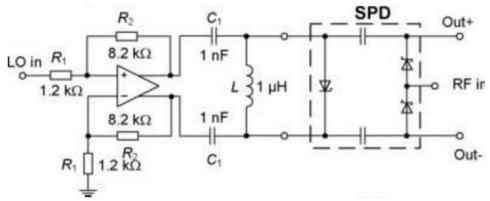

(c) UWB RX w/ Sub-Sampling [33]

Figure 4.2: **UWB RX Architectures.** A number of receive architectures have been proposed in research [49, 83, 95, 99], however all come with some level of drawbacks. Energy detection receivers suffer from high noise and the necessity of custom hardware. UWB reception with high-speed (gigasample) ADCs requires the use of expensive, power-hungry components or custom-built ICs with low narrowband noise suppression. Finally, UWB sub-sampling techniques come at the cost of high noise and an inability to suppress the effects of narrowband interference.

Traditional ultra-wideband receive architectures (shown in Figure 4.2) are designed to listen to a large amount of instantaneous RF bandwidth, yet UWB transmissions are relatively weak compared to the transmissions of their narrowband counterparts. If narrowband transmitters are in close proximity to one of these UWB receivers *and* their transmissions fall within the monitored bandwidth of the UWB receiver, disruption of the received channel impulse response will occur [5, 58].

Figure 4.3 shows the reconstructed channel impulse response under both normal conditions along with the presence of a narrowband interferer (1 mW EIRP) occupying the same channel.

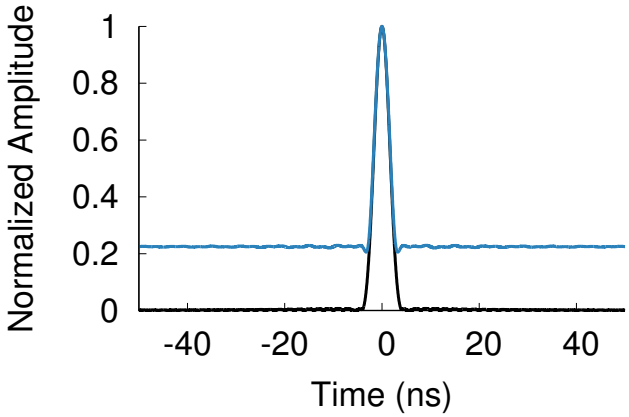

Figure 4.3: **Effects of Narrowband Interference on UWB CIR Reconstruction.** Reconstructed single-path UWB channel impulse response (black) with overlaid CIR in the presence of narrowband interference (blue). Simulation includes a narrowband transmitter operating with a 0 dBm output power placed at the same distance as the UWB transmitter. Any attenuation of the line-of-sight path may result in failure of detection in the presence of narrowband interference.

4.2.3 Decreasing Instantaneous Bandwidth Reduces Narrowband Susceptibility

Narrowband receivers are able to avoid interference from the presence of adjacent narrowband transmitters by either utilizing a bandpass filter before downconversion, at an IF frequency, or a low-pass filter at baseband. We leverage this fact to build a UWB receiver which is able to ignore those bands which are experiencing narrowband interference without the cost of complete failure such is the case with traditional UWB receive architectures.

The Harmonium architecture mirrors a traditional superheterodyne narrowband receiver frontend with the addition of a controllable, wideband local oscillator. The frequency-swept local oscillator allows for the reconstruction of the CIR in the frequency domain through the successive recordings of narrowband portions of the *channel frequency response (CFR)*. Once the CFR is known, its time domain dual (the CIR) can be obtained through an inverse discrete Fourier transform. Any Fourier coefficients which are known to be affected by narrowband interference can be discarded and estimated through the use of compressive sensing techniques [9]. Figure 2.7 describes these two different techniques for obtaining the CIR.

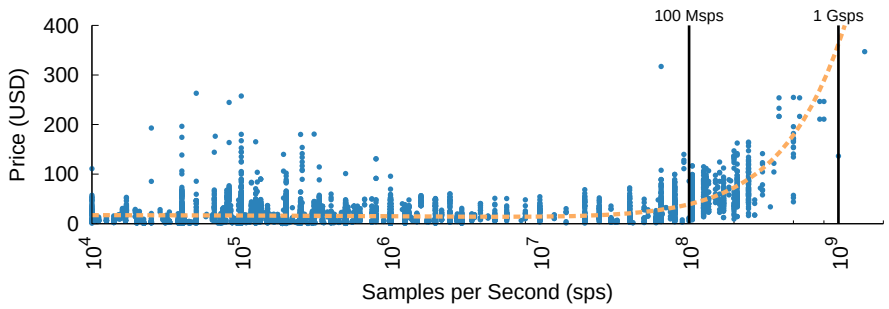

(a) Price vs. Sampling Rate

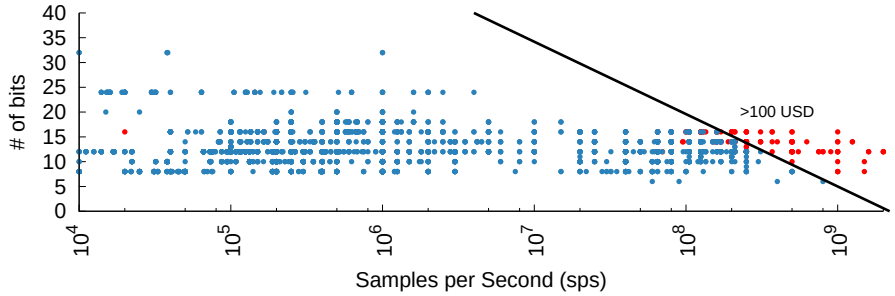

(b) Dynamic Range vs. Sampling Rate

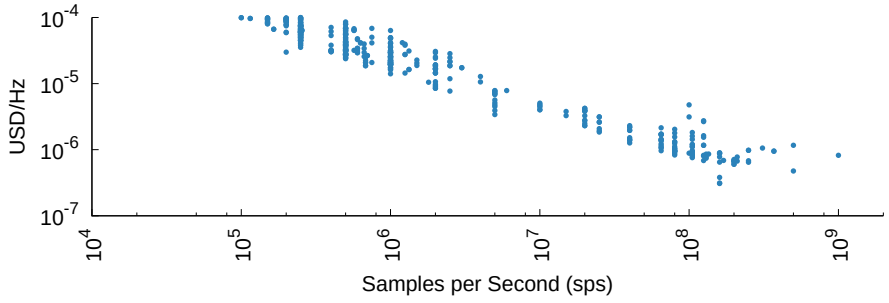

(c) Price/Sampling Rate vs. Sample Rate (16-bit ADCs)

Figure 4.4: **ADC Price Versus Sampling Speed.** We scrape all 9,716 of the available ADCs from the well known parts supplier DigiKey and collect the price per unit for the best bulk rate, discarding products only available in very small quantities (those with no bulk option available). In (a), it is shown that there exists a super-linear relationship between price and sampling rate above about 100 Msps, which is required for traditional UWB anchors. (b) shows that a realistic price break of 100 USD yields a tradeoff of sampling rate vs. dynamic range (interference rejection capability) at UWB bandwidths. In addition, (c) shows that the cost efficiency of ADCs does not currently improve significantly past a threshold of 100 megasamples per second, yielding an opportunity for cost-parity parallel bandstitching receiver implementations. Careful selection of ADC sampling rate and dynamic range is necessary for cost-effective anchor design.

4.3 Highly Constrained Applications Require Separate TX/RX Functionality

Another problem which arises through the use of SurePoint is that of a need for asymmetric approaches. For instance, as is the case with any RF transceiver, the receiver complexity is much greater than that of the transmitter. This imbalance could be leveraged to aid in creating low-power or cost-efficient tags for applications which are either size, weight, or power constrained. One example of an application which would benefit from separating the functionality of the transmitter and receiver are in that of SWaP-constrained quadrotors, like that shown in Figure 4.7.

To this end, a number of approaches have been shown for the generation of ultra-wideband signals [80]. UWB pulse generators have been designed using BJTs [28], boutique step recovery diodes [52], careful microstrip techniques [32], and custom ICs [89]. These approaches are difficult to implement due to their requirement of custom silicon, careful RF layout using expensive substrates, or employ electronic components which are expensive and difficult to source. Harmonium implements an ultra-wideband transmitter that achieves FCC-compliant UWB transmissions with commodity components and requiring no complex RF layout techniques.

4.3.1 The Transmit-Only UWB Tag

Harmonium tags produce high-bandwidth pulses using the step recovery effect of common RF bipolar junction transistors (BJTs) [28]. The step recovery effect creates a fast (sub-ns) state-change transition in semiconductor stack-ups due to the quantum effects of semiconductors recovering from a saturation condition. These fast state-change transitions (low- to hi-voltage or hi- to low-voltage) exhibit bandwidth exceeding 4 GHz, making them well-suited for generating UWB signals.

A crystal oscillator is first used to generate the stable pulse repetition frequency necessary for accurate channel impulse response characterization. This oscillator triggers a monoflop to generate a short-duration (multi-ns) driving signal for the NPN transistor. This short-duration pulse needs to be long and strong enough to drive the transistor into saturation. Once in saturation and the driving signal has transitioned back low, the BJT continues to bleed off charge until the observed step recovery effect takes place. The sharp transition from the conducting to non-conducting states cause a sharp rise in the output voltage present at the transistor collector. Finally, a DC-blocking capacitor and 9^{th}-order Chebyshev low-pass filter are used to attenuate unnecessary low-frequency components to

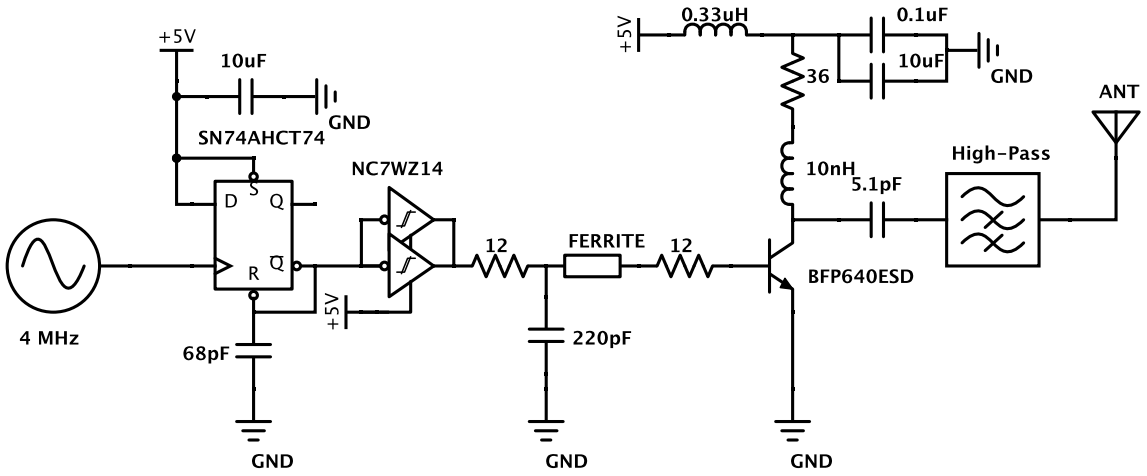

Figure 4.5: **Circuit Diagram for the Harmonium Transmit-Only Tag.** Shows the detailed interconnection between the oscillator, monoflop generator, and BJT transistor. Additional passives are necessary for FCC-compliant pulse shaping. Total tag cost in modest quantity is approximately $4.50 per tag.

generate FCC-compliant UWB signals. Figure 4.5 shows a schematic of our tag design.

Figure 4.8a shows the fabricated pulse generator tag. The tag PCB is constructed using a Rogers 4350 PCB laminate material. The tag is set to generate a pulse train at a 4 MHz repetition frequency using a crystal oscillator. This pulse repetition rate allows for channel delay spreads of up to 250 ns, which we experimentally determine to be adequate for many indoor environments. The fabricated tag occupies 1.5 cm^2, weighs 3 grams (weight is of key importance for quadrotor handling and power [64]), has an active power draw of 75 mW (approximately $1/4$ of SurePoint's 280 mW), and costs only $4.50 USD in modest volumes.

4.3.2 The Receive-Only Harmonium Anchor

Each anchor uses three UWB antennas [3] to receive tag transmissions and provide antenna diversity, which has been shown to improve ToA estimation performance in prior work [44]. An RF switch selects different antennas over successive localization measurements. The switched antenna then feeds an LNA and mixer circuit to enable bandstitching at each anchor. Each anchor mixer is fed from a central frequency-stepped local oscillator source to facilitate synchronous bandstitching across all four anchors. The local oscillator (LO) signal generation board uses an ADF4159 frequency synthesizer that controls the frequency of an RFVC1802 wideband VCO. The LO sweeps from 5.312 GHz to 4.32 GHz in 32 MHz steps. With an intermediate frequency of 990 MHz at each anchor, this in approximately 1 GHz of bandstitched bandwidth from 3.33 GHz to 4.322 GHz.

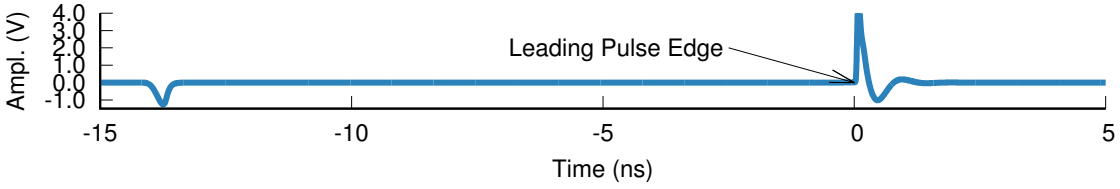

(a) Simulated Pulse

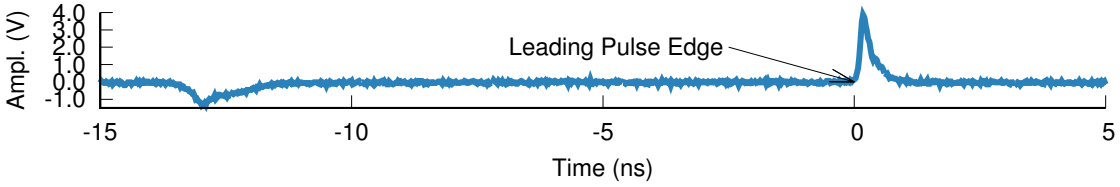

(b) Measured Pulse – No High Pass Filter

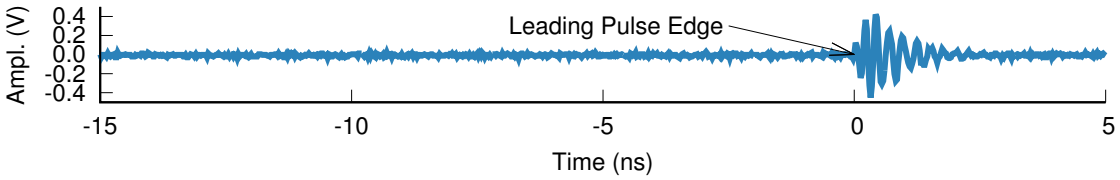

(c) Measured Pulse – With High Pass Filter

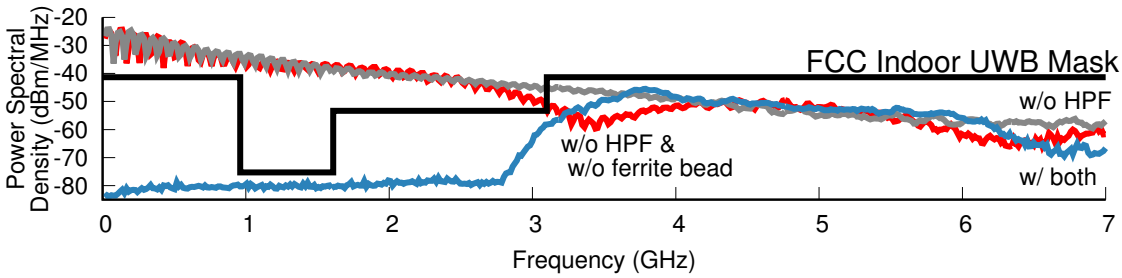

(d) Measured – Frequency Domain

Figure 4.6: **Harmonium Tag Performance.** Simulated and measured time- and frequency-domain characteristics of the Harmonium pulse generator. The generated pulse is 275 ps wide (FWHM) and occupies more than 7 GHz of bandwidth. The frequency content below 3.1 GHz is stronger than allowed by FCC UWB guidelines, requiring the use of a high-pass filter to attenuate low-frequency content.

Figure 4.7: **A Harmonium Tag on a Micro Quadrotor.** This device has a maximum payload of 10 g and has only a 170 mAh battery. The Harmonium tag mass (3 g), volume (2.3 cm^3), and power requirements (75 mW) are low enough to provide high-fidelity localization even for this constrained device. This enables real-time (19 Hz) location updates from a flying micro quadrotor.

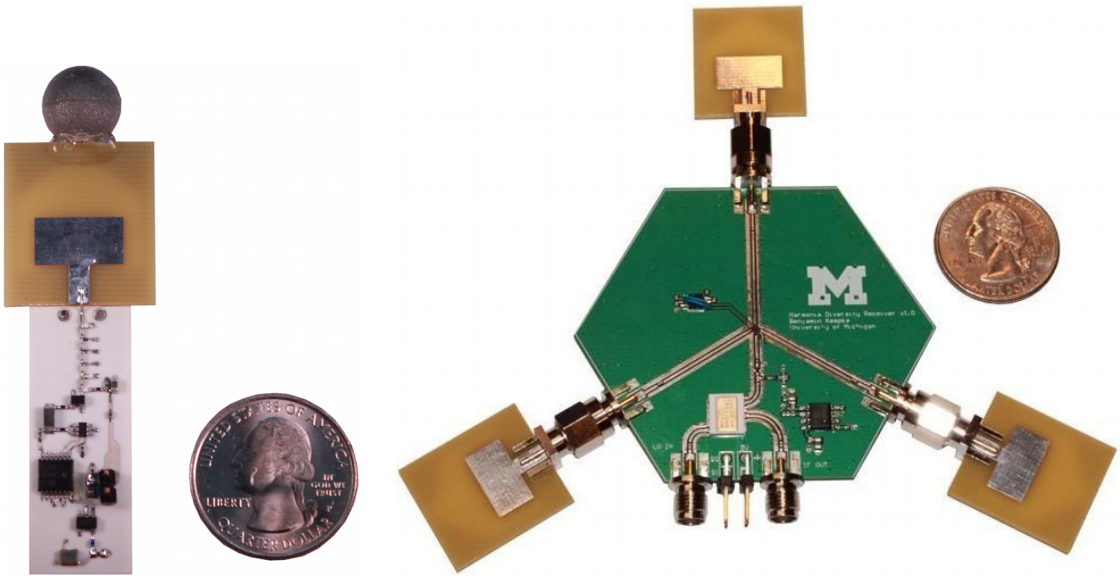

(a) Tag (b) Anchor

Figure 4.8: **Harmonium tag and anchors.** Tags measure 2.2 cm x 6.3 cm, contain pulse generation circuitry, and are printed on Rogers 4350 PCB substrate. Anchors consist of a centralized local oscillator (LO) frequency generator and separate RF front-ends for down-converting and digitizing the received pulse signals measured at each anchor.

The resulting mixed intermediate frequency (IF) signal from each anchor returns to a dedicated USRP1 [23] front-end for final down-conversion, digitization, filtering, and data transport to an attached PC. The DBSRX2 daughterboard first converts the 990 MHz IF signal to baseband for ADC sampling. The USRP1 uses 64 Msps baseband ADCs with a bit depth of 12 bits for each of the baseband quadrature channels (in-phase and quadrature-phase). Since the USRP1 uses USB 2.0 to transfer baseband data to the host PC, the resulting 64×10^6 Msps $\times$ 12 bits/sample $\times$ 2 channels = 1536 Mb/sec of baseband data is too much to pass unprocessed to the host PC. Instead, the raw baseband data is comb filtered and decimated to decrease the overall bandwidth required of the host PC data interface.

All signal processing and LO interfacing logic is implemented using a custom FPGA image loaded onto the USRP1's Spartan 3 FPGA. The system repeatedly sweeps the entire bandwidth sequentially across all three antennas, producing localization estimates at 19 Hz. The anchors used in this evaluation cost approximately $750 each due to the high cost of COTS SDRs, yet these could conceptually be replaced with a custom SDR implementation to significantly reduce anchor cost. This, coupled with the advent of inexpensive, integrated wideband synthesizer/mixer RFICs such as the RFMD RFFC5072 could reduce anchor cost to $100.

4.3.3 The Narrowband Retrofit

An alternative to the techniques described above – if available – is to modify the transmission path of a currently-existing RF frontend. We have demonstrated this capability in prior work [43] to modify the transmission path of a typical narrowband radio to support UWB transmission. Depending on the application, either one of these two methods can be sufficient to achieve UWB-level localization accuracy in heavy multipath environments.

4.3.4 Localizing Harmonium Tags

From the base UWB receiver architecture, Harmonium obtains the precise time that pulses are received at each anchor according to the anchor's frame of reference. For this reason, Harmonium anchors all share a tightly synchronized, global time reference. Anchors calculate the offset between the arrival of the tag's transmitted pulse and the global clock pulse.

Calculating position based on time-difference-of-arrival (TDoA) is well-known as multilateration, and uses the same principles used by GPS to perform geolocation. Algorithms have been developed to determine location in three dimensions based on the addition of one or more time-of-arrival estimates [7, 8, 10, 31, 79]. To minimize the effect of antenna

cross-polarization, which can substantially attenuate a signal, Harmonium uses three antennas at each anchor and computes the tag's position using the earliest time-of-arrival at each anchor.

4.4 Lowering the Resistance to UWB CIR Adoption

In previous chapters we have outlined the need for UWB generation and measurement, a modular system built from commodity components to realize a localization system using UWB, and the main drawbacks of the system as it stands, including the lack of application coverage. In this chapter, we have outlined methods to circumvent these issues along with a real-world evaluation of these applications. The next chapter will shift focus to protocol-level design modifications to aid in the realization of multi-tag SurePoint systems.

4.5 Harmonium Performance Assessment

We evaluate the Harmonium prototype on precision, accuracy, consistency, and system burden—weight, volume, and power requirements. We conduct all experiments in an approximately rectangular $4.6 \times 7.2 \times 2.7$ m room in a commercial building with heavy multipath characteristics. We assign the origin to a floor-level corner and coordinate axes run along each of the orthogonal wall edges. We install a NaturalPoint OptiTrack motion capture system [65], calibrated to a sub-mm accuracy, in the room to provide ground truth measurements for all experiments. Harmonium achieves a median 14 cm error with a 90^{th}-percentile error of 31 cm and median precision of 9 cm while drawing only 75 mW with a 3 g tag.

4.5.1 Stationary Precision (Figure 4.10)

We place a tag at fifteen fixed positions in space, taking roughly forty samples at each position, to measure the typical magnitude of position estimation noise from system and environmental noise. Figures 4.10a and 4.10b show ground truth locations and point cloud estimates for each position in line-of-sight (LOS) conditions. Harmonium achieves 14 cm median error with 9 cm median precision.

For the final stationary experiment, shown in Figure 4.10d, we introduce a strong narrowband interferer by radiating a modulated 3.6 GHz signal with a nearby USRP. While the overall median error and precision, 28 cm and 17 cm respectively, continue to perform well, certain physical spaces fail completely, such as position #14 which exhibits 217 cm median

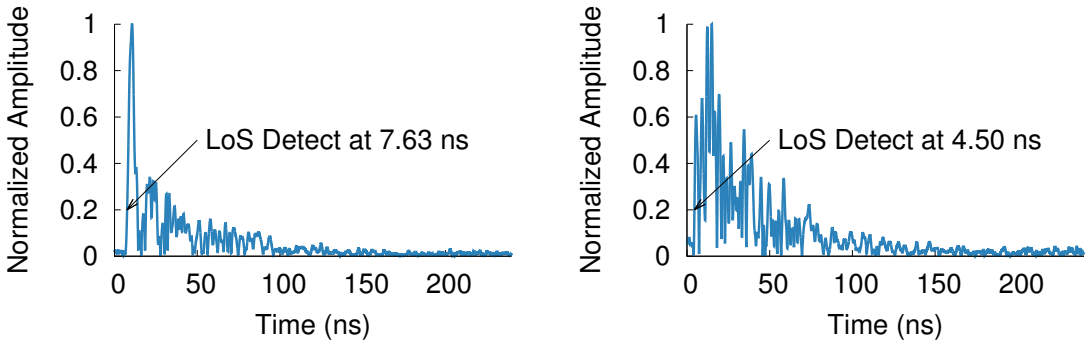

(a) Anchor 1

(b) Anchor 2

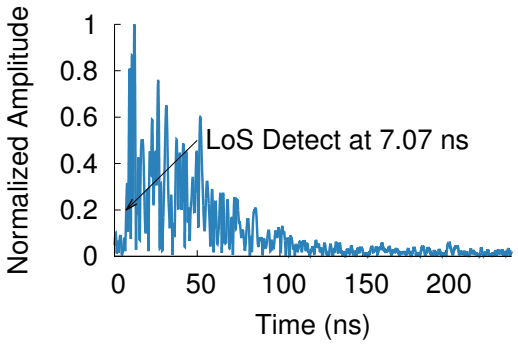

(c) Anchor 3

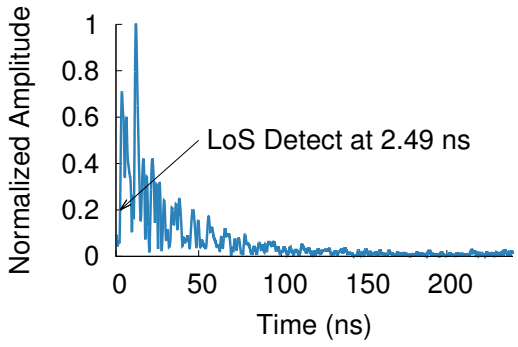

(d) Anchor 4

Figure 4.9: **Example CIR Measured Using the Harmonium RX Architecture.** The time-of-arrival for the CIR's leading edge is used as an estimate for the arrival time for the line-of-sight path. Accurately determining the LoS arrival time is the key to determining tag position with low error.

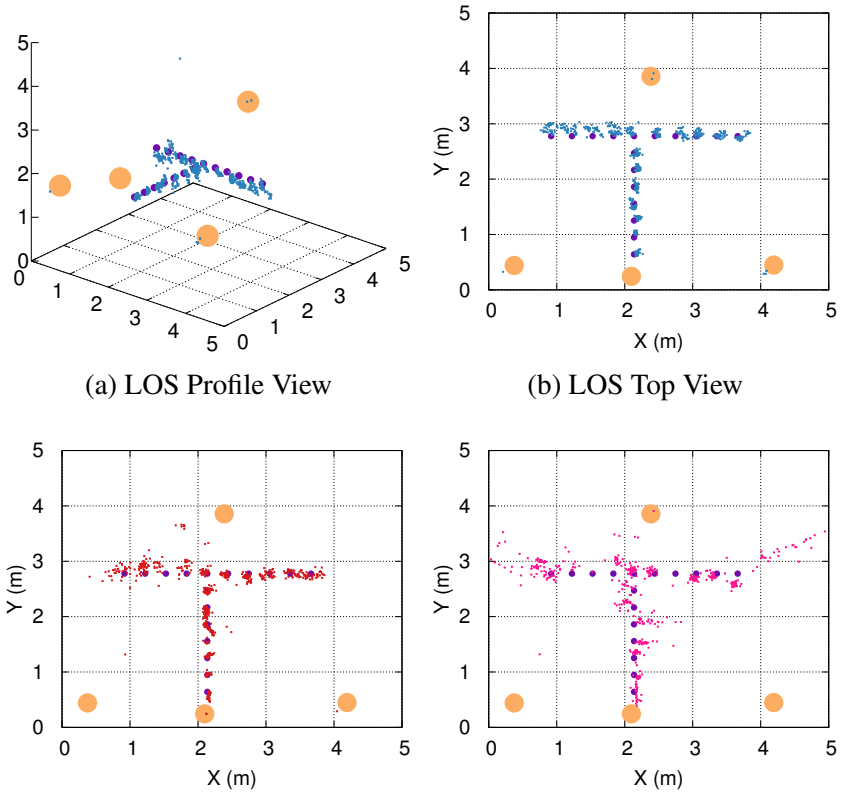

(a) LOS Profile View

(b) LOS Top View

(c) Through-Wall Top View

(d) NB Interference Top View

Experiment	Median Precision	90th %-ile Precision	95th %-ile Precision	Median Accuracy	90th %-ile Accuracy	95th %-ile Accuracy
LOS	9 cm	16 cm	37 cm	14 cm	31 cm	37 cm
Through-Wall	13 cm	38 cm	51 cm	16 cm	42 cm	53 cm
Interference	17 cm	43 cm	107 cm	28 cm	136 cm	201 cm

(e) Key metrics

Figure 4.10: **Harmonium Static Position Performance.** We place Harmonium at fifteen known locations and capture roughly 40 position estimates at each point. First we capture the line-of-sight (LOS) base case. Then we evaluate through-wall performance by occluding the anchors with drywall. Finally, we introduce a narrowband interferer strong enough to completely knock out a commercial UWB system and observe Harmonium's performance. Harmonium exhibits minor (2 cm) performance degradation in the through-wall case and only $2\times$ loss in median accuracy in the face of strong narrowband interference, demonstrating the efficacy of Harmonium's bandstitching architecture.

error with 38 cm precision. Recently, the first commercially-accessible UWB transceiver, the DecaWave DW1000 [18] was released. While building and evaluating a complete localization system using DecaWave to compare against is beyond the scope of this paper, we do validate one of our previous claims that motivated the bandstitching-based approach, and find that a pair of DecaWave chips fail to communicate in the presence of the same narrowband interferer.

These experiments give a sense of the consistency of position estimates obtained with Harmonium. Due to the approximately normal distribution of position estimation noise across each dimension, a reduced variance in position estimation noise can be obtained by taking a moving average of position estimates. While this will decrease the average position error, it has a cost of reduced position update rate. All following experiments are performed using raw position estimates without any temporal filtering of the data.

4.5.2 Quadrotor Flight Path Reconstruction (Figure 4.11)

We next evaluate Harmonium in a motivating application domain: real-time tracking of indoor quadrotors. The Crazyflie Nano is a 19 g, $9 \times 9 \times 2$ cm quadrotor with a 170 mAh battery and a payload capacity of only 10 g [4]. The existing motors and electronics draw approximately 1400 mA (5180 mW at 3.7 V) while hovering, so the Harmonium tag power draw only reduces flight time by 1.4%. The additional weight dominates the additional power draw required for the Crazyflie to maintain hover. With an approximate 200 mW/g of additional payload, the quadrotor would require an extra 600 mW of power to maintain hover with an affixed Harmonium tag.

We affix a Harmonium tag and fly the quadrotor around the indoor space. Figure 4.11 captures a trace of this flight. The flight exhibits a median error of 14 cm and 90$^{\text{th}}$ percentile error of 35 cm. Empirically, significant errors are clustered in space and time, suggesting that there is a physical root cause and that temporal filtering will be insufficient to resolve the errors. We explore this observation further in the next experiment.

4.5.3 Consistency on a Static Path (Figure 4.12)

While the quadrotor experiment demonstrates Harmonium's ability to reconstruct a challenging arbitrary path, we are also interested in the reproducibility of Harmonium's position estimates over time. In Figure 4.12, we place a tag on a model train and record ten laps around the fixed track. During this experiment, we move about the space normally, perturbing the multipath environment between samples at the same point in space. Figure 4.12c shows an aggregate point cloud of all ten laps and the variation across laps. While the

position error varies around the track, the variance is consistent at each location, that is the standard deviation of position error is relatively constant. This suggests that the measurement error has a physical root cause based on the properties of specific points in the space.

4.5.4 Pulse Generation and Regulatory Compliance (Figure 4.6)

As pulse generation quality leads directly to spectral usage, which in turn informs location quality, Figure 4.6 evaluates the expected and actual performance of Harmonium's pulse generation circuitry. The addition of the high-pass filter, necessary for regulatory compliance, abbreviates the tag's effective bandwidth. However, the design is still able to achieve nearly 3.5 GHz of bandwidth.

4.5.5 System Microbenchmarks

The Harmonium design expressly introduces an asymmetry between tags and anchors to minimize the burden of introducing Harmonium tags to devices to be localized. Here we quantify how burdensome the realized tag design is and the impact on the Harmonium anchor. The Harmonium tag is made of a 3.9×1.5 cm PCB with a 2.4×2.2 cm UWB antenna. The whole tag fits within a $3.9 \times 2.2 \times 0.2$ cm bounding box, or about 1.5 cm^3. The tag weighs only 3 g and draws only 75 mW. At a 19 Hz update rate, the tag uses 3.9 mJ per location estimate. The Harmonium anchors consist of a central 6.7×5.8 cm PCB with three 2.4×2.2 cm UWB antennas mounted co-planar at 120° offsets. One USRP1 can service up to two Harmonium anchors. The data from one USRP1 (two anchors) nearly saturates a USB 2.0 bus, requiring USB 3.0, more than two bus controllers, or multiple machines to support more than four Harmonium anchors. One 3.2 GHz Xeon core can solve a position estimate in 231 ms. At least five parallel cores are required to keep up with Harmonium's 19 Hz update rate.

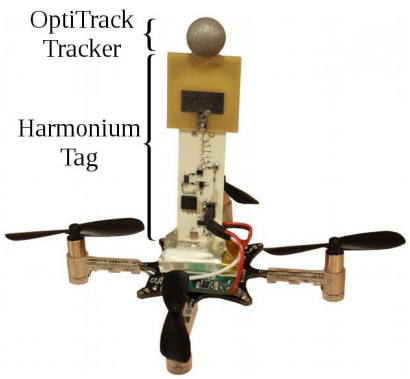

(a) Quadrotor with Tag

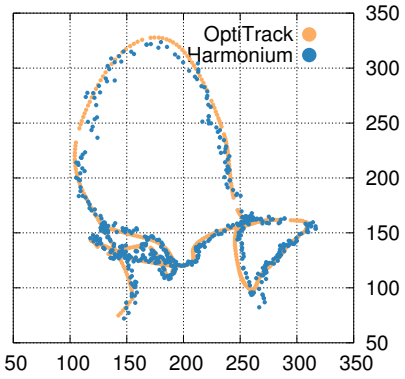

(b) Top View

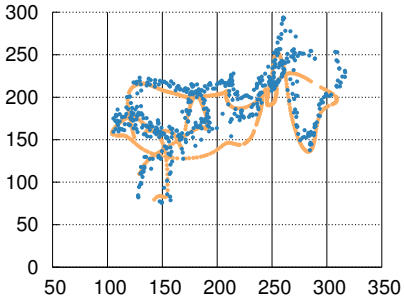

(c) Side View from X

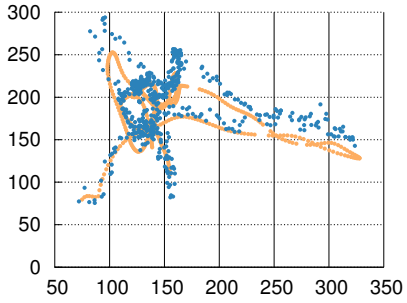

(d) Side View from Y

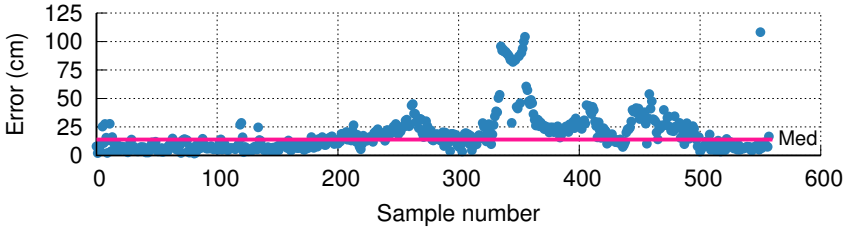

(e) Error

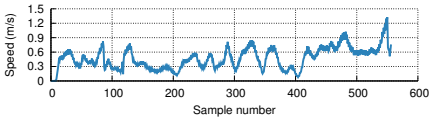

(f) Speed

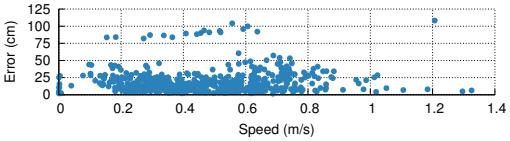

(g) No Correlation Between Speed and Error

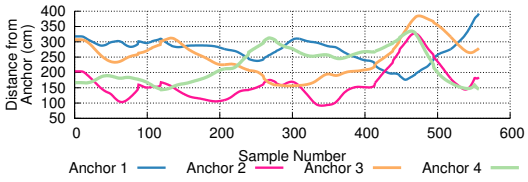

(h) Tag to Anchor Distance

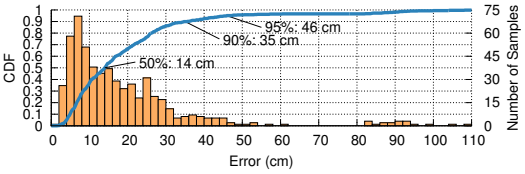

(i) Error CDF and Histogram

Figure 4.11: **Harmonium Quadrotor Tracking Performance.** Point-cloud of location estimates and CDF of location error tracking a quadrotor. Harmonium shows no increase in error up to the 1.4 m/s top speed of the quadrotor, nor does Harmonium severely burden the quadrotor's ability to fly, adding less than 15% to the mass.

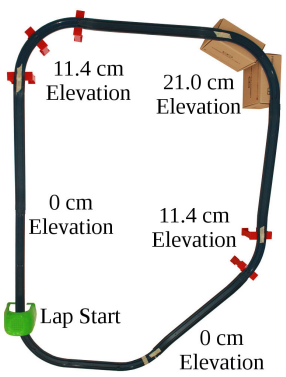

(a) Train Setup

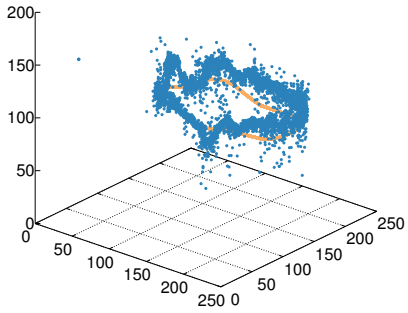

(b) Profile View

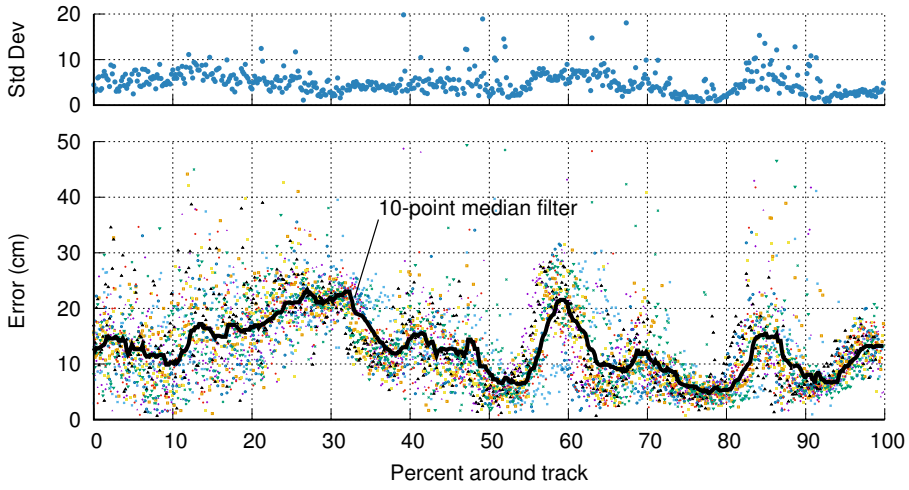

(c) Consistency of Position Error Across Laps

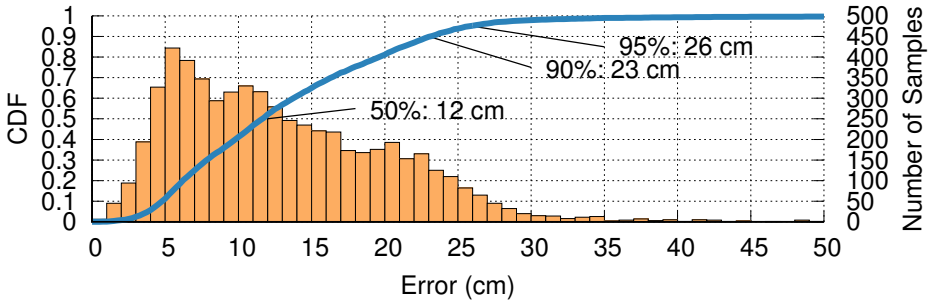

(d) Error CDF and Histogram

Figure 4.12: **Harmonium Spatial Error Correlation.** Point-cloud of location estimates and CDF of errors tracking ten laps of a model train around a track. Errors are consistent in space, implying a physical root cause.

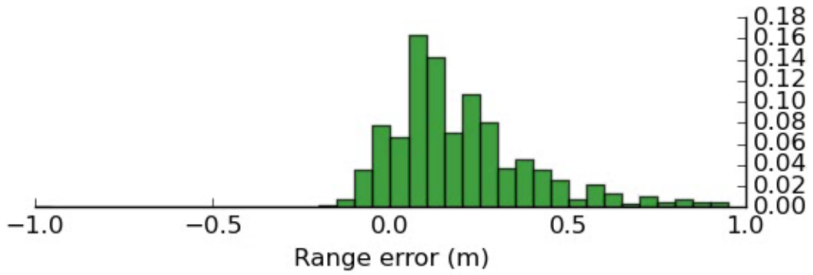

(a) Distribution of Range Errors from Channel Model

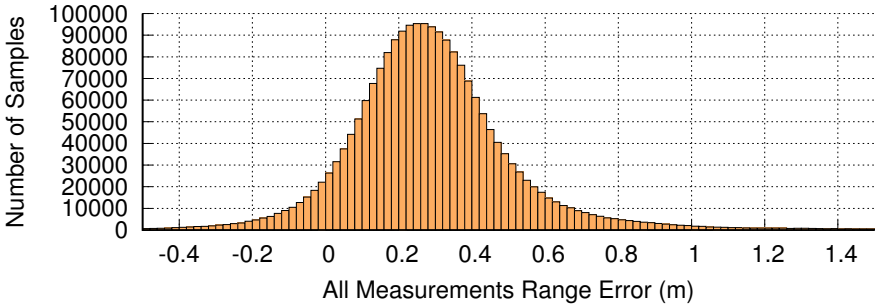

(b) Distribution of Range Errors from Measurements

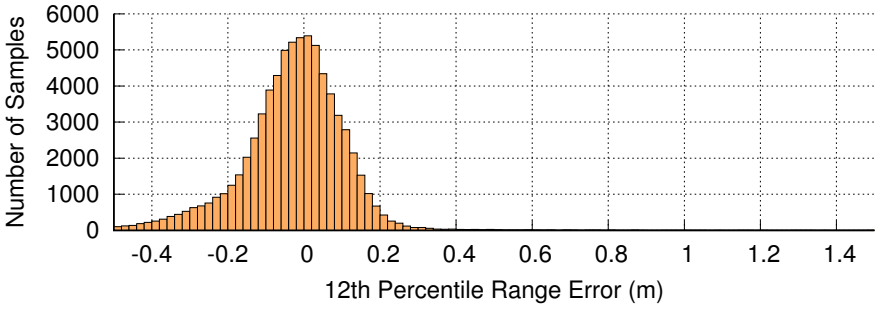

(c) Range Errors after Diversity Aggregation

Figure 4.13: **Ranging in the UWB Channel.** While UWB provides higher fidelity range estimates than narrowband radios, the estimate is still imperfect. Using an 802.15.4a channel model from Molisch et al. [62], take 1000 samples of an active multipath environment and find that the estimated range between two nodes can vary by over a meter. We then collect samples in a real world environment and find that it follows the expected distribution. From these analyses, assuming an unlimited number of independent observations, the 12th percentile will yield zero range error. The final plot captures the 12th percentile estimates from the real-world samples, reducing the median error from 0.27 m to 0.12 m and 95th percentile errors from 1.31 m to 0.35 m.

4.6 Summary

In this chapter, we have described the design of two separate pieces of hardware which are important to the realization of ultra-wideband localization technologies. First, we have introduced a UWB transmitter design which is based solely on commercially-available components and consumes order-milliwatts to generate ultra-wideband signals. Secondly, a bandstitched ultra-wideband receiver has been introduced which allows for the measurement of the indoor ultra-wideband channel impulse response. Together, these two building blocks synthesize a system which is able to track fast-moving objects with minimal additional cost or complexity at the tag being tracked.

The bandstitched ultra-wideband receiver provides a powerful and configurable diagnostic tool which can allow for a much deeper analysis of the channel impulse response beyond that currently available with commercial ultra-wideband transceiver topologies. In the next section, we will describe a specific analysis of the channel impulse response which can aid the tracking of ultra low-power tags which are unable to afford the power consumption requirements of the Harmonium tag, yet have the specific characteristic of remaining stationary for long periods of time, allowing for us to trade off localization update rate for tag power.

CHAPTER 5

Slocalization: A Batteryless Backscatter UWB Transmitter for Asset Tracking

Recent advancements in ultra low-power electronics have given rise to a new class of devices which are able to operate using only the energy which can be harvested from the surrounding environment. A variety of different methods have been demonstrated to harvest energy through vibration, light, radio, and a variety of different sources. The energy attainable through these methods is usually quite small, on the order of a few *micro-watts*. Due to these minimal average power requirements, energy-harvesting devices must either accommodate the restricted power requirements or duty-cycle their operation.

The main advantage of energy harvesting devices lies in their ability to operate independently with little to no ongoing maintenance due to their ability to operate without batteries. This makes the energy harvesting architectures particularly well-suited to the application of asset tracking.

Asset tracking is commonly used to determine the location of objects which stay motionless for extended periods of time. Motivating examples of asset tracking include finding a lost set of keys, locating a book on a shelf, or tracking inventory in a warehouse. Due to the large number of potential tagged objects, battery maintenance is undesired and often implausible.

The following section describes the design and evaluation of a tag which is able to achieve accurate localization operations using less than one microwatt of active power. Instead of utilizing an ultra-wideband transmitter to achieve a high-resolution measurement of the RF channel environment, the Slocalization tag achieves high-accuracy ranging measurements by reflecting the the transmissions emanating from anchors situated throughout the indoor environment. This design trades update rate for tag power – a tradeoff which is reasonable for the domain of asset tracking where objects are expected to remain stationary for extended periods of time.

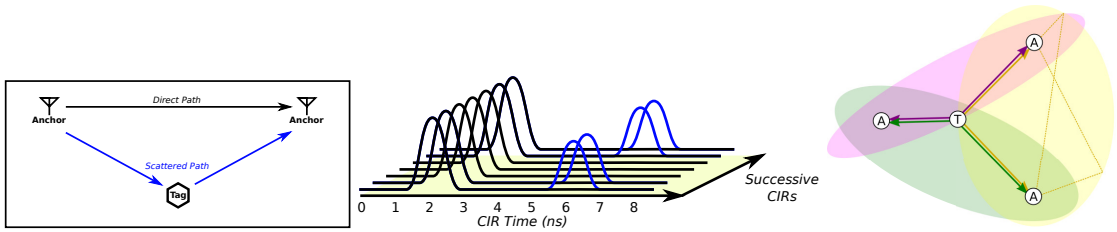

(a) Backscatter Setup (b) CIR Behavior Over Time (c) Location from Time

Figure 5.1: **Slocalization Concept of Operation.** Anchors with known positions emit periodic pulses that sound the ultra wideband channel. A Slocalization tag modulates its antenna to either reflect or absorb this signal, perturbing the channel impulse response (CIR) over time. Anchors use the time difference of arrival between the direct path between anchors and the backscatter path reflected from the tag to form ellipsoids of possible tag locations. The intersection of sufficient ellipsoids yields the absolute position of the tag.

5.1 Tags that last forever: An emphasis on tag power

Previous ultra wideband localization systems are an especially interesting class to consider in the context of "slow" localization. The quintessential sensor networking technique for reducing energy consumption is to find a way to reduce the duty cycle of a node. If the argument is truly that devices rarely or never move, then perhaps running traditional localization systems at millihertz duty cycles is the right approach.

One immediate drawback for such a design is poor peak to average power ratio, a prohibitive design point for battery-based systems. The capacitive storage banks of energy harvesting architectures, however, are well suited to intermittent high current operation. Furthermore, the infinite lifetime provided by energy harvesting is very well aligned to the eventual vision of localizing all physical things. High peak power requirements do still require sufficient storage (in terms of capacitor volume and board area) to support operations, however. To quantify these trade-offs, we look at the state of the art in low power decimeter-accurate localization systems. For such a design, we only consider systems where the underlying localization mechanism can achieve an immediate fix.

5.1.1 With Commercial Transceivers

The lowest power decimeter-accurate single-fix localization with traditional radios is SurePoint, with 80 ms long ranging events at 280 mW, or 22.4 $^{mJ}/_{range}$ [42]. SurePoint includes additional overhead beyond the ranging events for scheduling and maintaining time slots, however for sake of argument let us assume that the very low duty cycle will effectively eliminate interference and that there is zero static power draw between range events. To

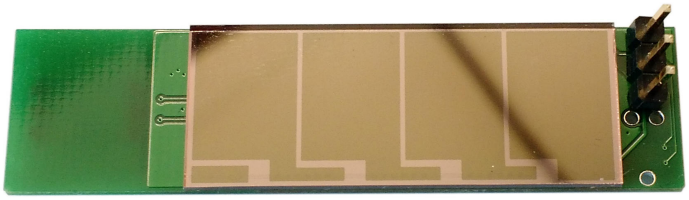

(a) Front

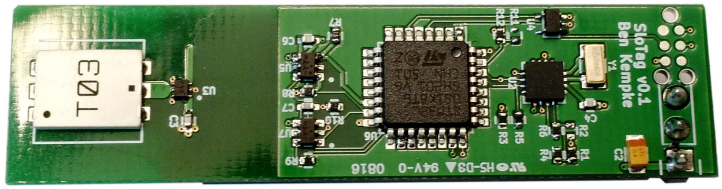

(b) Back

Figure 5.2: **Realization of the Slocalization UWB Backscatter Tag.** Utilizing only commodity components, the Slocalization backscatter tag operates on only 1 µW of power, enabling its use in energy harvesting applications. The tag is powered with a small (35 mm × 14 mm) solar cell and limited energy storage (47 µF) to demonstrate its applicability to demanding energy harvesting applications.

realize Slocalization's 1 µW, SurePoint could only range once every 6.2 hours.

For energy harvesting applications, SurePoint's 3.3 V operating level raises additional concerns. Using the harvesting and activation circuit from Monjolo [16], whose regulator is generously 80% efficient across the 2-0.35 V input 3.3 V and 100-200 mA output range, requires 28 mJ in the storage capacitors, or roughly 14 cm^2 of board area for similar capacitors.

The primary energy cost in SurePoint is the 145 mA DecaWave ultra wideband transceiver. Even an order of magnitude improvement in transceiver energy could still only realize one transmission every 40 minutes at 1 µW.

5.1.2 With Impulse Frontends

Prior systems have also identified the transceiver as the most costly component, and replace it with a simpler and cheaper ultra wideband pulse generator. The current lowest power decimeter-accurate single-fix localization system is Harmonium [45]. Ignoring the issue of synchronizing tag ranging pulses with Harmonium's bandstitching, capturing a location fix requires the tag to transmit for 53 ms at 75 mW, or 4 $^{mJ}/_{range}$. To realize a 1 µW average power budget, a Harmonium tag could transmit ranging pulses every 1.1 hours.

The Harmonium impulse generation circuit relies on exploiting the step recovery ef-

fect in RF BJTs. This requires the tag to have a relatively high operating voltage of 5 V. Again considering the Monjolo energy harvesting frontend, reaching 5 V adds an additional burden for energy harvesting designs. For a 5 V, 15 mA output, the regulator efficiency improves to 85% thus requiring 4.7 mJ in the storage capacitors, or roughly 2.4 cm^2 of board area for storage capacitors alone.

5.1.3 Backscatter Required

We suspect that some of this domain space exploration does not come as a surprise to backscatter experts. The simple reality is that the energy required to open and close a switch (to reflect RF energy) is so much less than the energy required to radiate RF energy that even with a five order of magnitude increase in "transmission duration", the backscatter approach consumes significantly less tag energy for a single location fix.

5.2 Slocalization Compared to RFID

While the backscattering concept dates back decades [69, 85], there has been a recent resurgence in research around backscatter, extending the concept from beaconing simple identifiers to high bandwidth communication [77, 98], highly parallel communication [37], leveraging ambient environmental signals instead of active interrogators [57], or even motion capture [87, 96]. Localization is a similarly mature line of research, however, with the advent of new FCC regulations in 2002, the last decade has seen an explosion of interest in ultra wideband for localization due to the greatly improved resolution it can provide indoors [14, 30, 59, 60].

In traditional backscatter systems, an interrogator (e.g. an RFID reader) emits a powerful, well-known signal – often a pure sine tone. Tags in the environment modulate the impedance of their antenna by opening and closing a switch, changing their antenna from being highly reflective to highly absorptive. A receiver captures these reflections and uses them to recover data from the tag. The key insight in backscatter is that it enables a vast energy asymmetry between the anchor (interrogator) and the tag, as the energy cost of actuating a switch to change impedance is very low.

Some prior work has proposed the viability of ultra wideband backscatter in theoretical analyses [19, 34]. As theoretical systems, these designs rely heavily on antenna and channel models to validate system design choices. Unfortunately, the standard 802.15.4a channel model [63] is not well suited to modeling a "two-way" signal, i.e. a backscatter reflection, requiring simulations to mix in motion models or employ statistical tricks to try to model

a complex, indoor ultra wideband backscatter channel [34]. D'Errico et al. further explore how to design a hybrid system that includes a conventional RFID frontend for wakeup and energy harvesting, however as previously stated our Slocalization design is independent of energy frontends and should be amenable to such a hybrid design.

5.3 Tag Architecture: UWB Backscatter Transmission

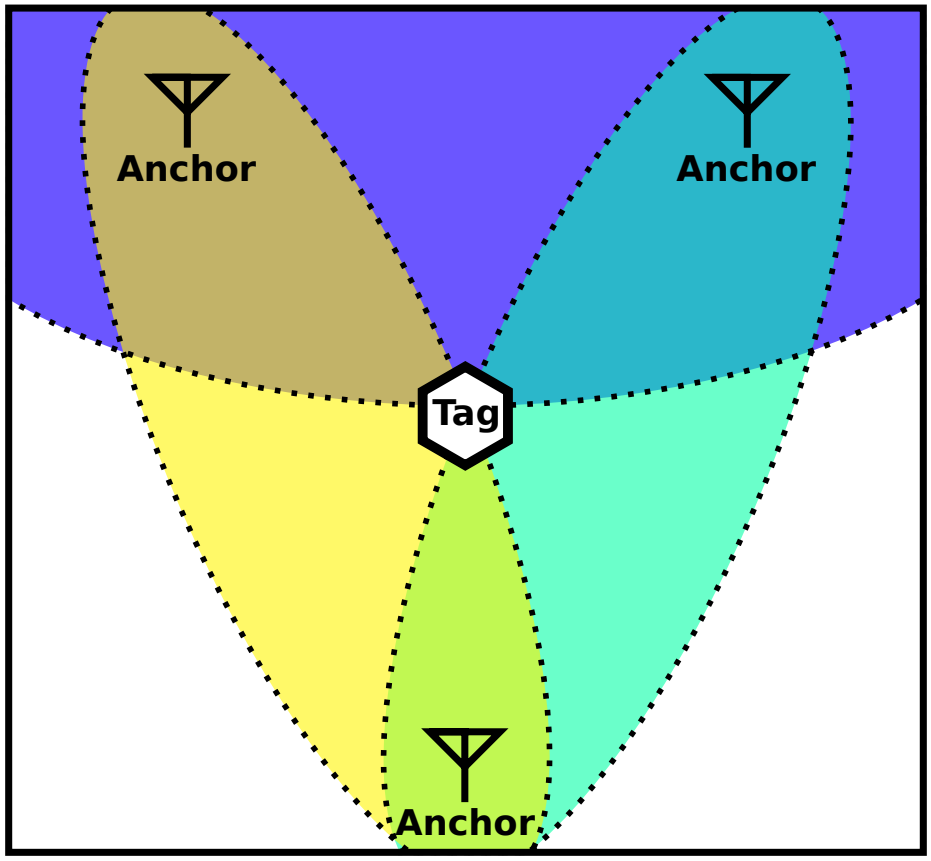

Figure 5.3: **Logical Organization of the Slocalization Backscatter Localization System.** Tags are localized by measuring perturbations of the channel impulse response observed between anchors situated throughout the environment. Long integration times are used to increase SNR to a level acceptable to observe these minute variations. The leading edge of the tag's CIR corresponds to the time difference-of-arrival between the direct path between anchors and the backscatter path traversed through the tag. Once these TDoAs are known, the tags are localized using ellipse-based TDoA localization techniques.

Backscatter techniques have seen widespread use in systems targeting ultra low-power tag operation. Unlike the active tags discussed so far, backscatter transmitters only transmit information through the modulation of the tag's antenna impedance. This acts to *change*

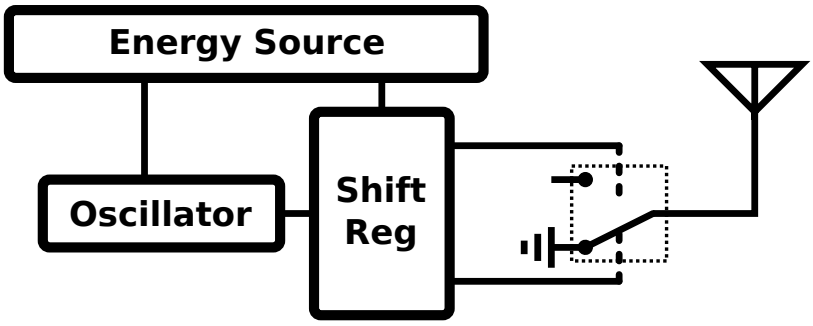

Figure 5.4: **UWB Backscatter Tag Design.** In order to provide backscatter communication, a UWB antenna and RF switch are used in conjunction to modulate the reflective characteristics of the RF channel. A shift register stores a PN code for the tag to emit. A high stability oscillator clocks the shift register to drive backscattered communication.

the multipath environment by electrically controlling the presence and/or phase of the tag's reflection.

Backscatter techniques have been used extensively in narrowband applications such as RFID (RF energy harvesting tags) and even with narrowband localization systems such as WiDraw [76]. However, their accuracy in stationary RF environments is again limited by the small amount of utilized bandwidth.

The same techniques can be applied in the ultra-wideband case. Figure 5.4 shows the architecture of the "Slocalization" backscatter tag. The impedance of an ultra-wideband antenna is modulated between a direct short and open through the use of an RF switch. The state of the switch is dictated by a microcontroller which can choose the pattern in which the multipath environment is modulated.

The Slocalization tag switches the antenna state once every second, modulating it with a pseudo-random code to provide code separation between unique tags. The UWB anchor (not yet designed) will then periodically sound the channel to calculate the observed channel impulse response. Any change in the Slocalization antenna state will result in a deviation in the channel impulse response with an offset corresponding to twice the difference in distance between the anchor's transmitter and receiver.

Successive channel impulse response measurements will be analyzed and averaged across time to increase SNR. According to simulation, the amount of time required to achieve a SNR of 20 dB (the minimum required SNR to perform leading edge detection) versus the tag position is shown in Figure 5.6. This will dictate the minimum anchor density requirements for asset tracking systems depending on the required update rate of asset position.

5.4 Low Backscatter Power = Long Integration Times

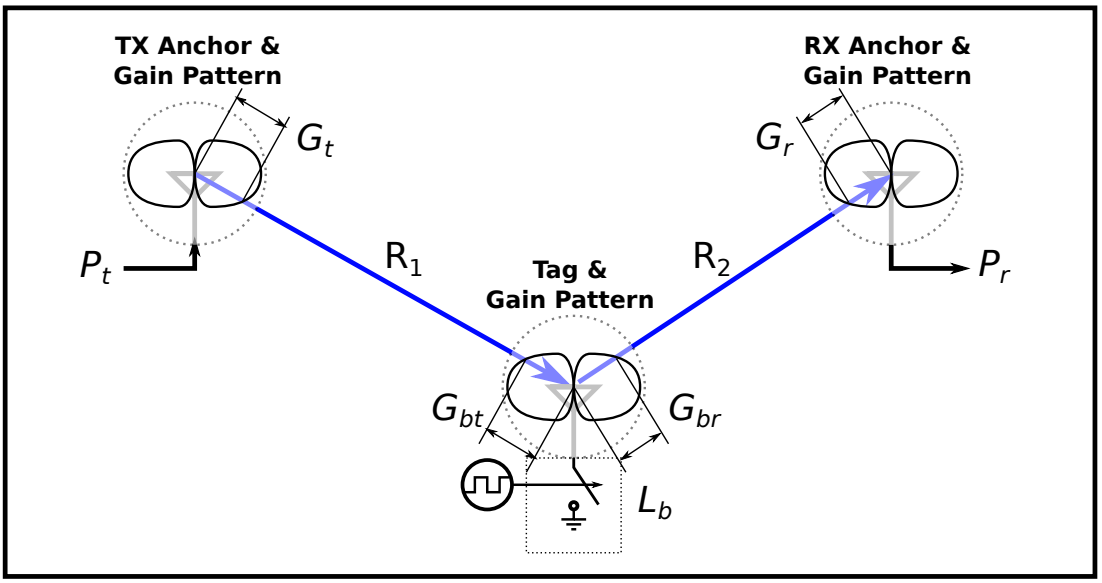

Figure 5.5: **Slocalization Link Budget.** As the backscatter tag is not an active transmitter, its localization relies on the measurement of reflected signals from another active transmitting source. A number of losses between transmitter and receiver are accrued due to path loss to the tag from the transmitter, path loss to the receiver, and losses internal to the tag. Slocalization requires long integration times to ameliorate these effects.

Backscattered signals are much weaker than those from an active transmitter as they must travel twice the distance. Recovering backscattered UWB signals is further confounded by limitations on UWB transmission power [25, 35]. The link budget for a Slocalization tag consists of three parts, also shown visually in Figure 5.5:

1. Path loss from transmitter to tag
2. Loss at the Slocalization tag
3. Path loss from tag to receiver

The total combined path loss can be summarized through an adaptation of the Friis transmission equation:

$$\begin{aligned} P_r = \; & P_t + G_t + G_{bt} + G_{br} + G_r + \\ & 20 \times \log_{10}\left(\frac{\lambda}{4\pi R_1}\right) + 20 \times \log_{10}\left(\frac{\lambda}{4\pi R_2}\right) - L_b \end{aligned} \tag{5.1}$$

where P_t and P_r are the transmitted and received power, G_t and G_r are the anchor's transmit and receive gains, G_{bt} and G_{br} are the gains of the tag's antenna from the perspective of

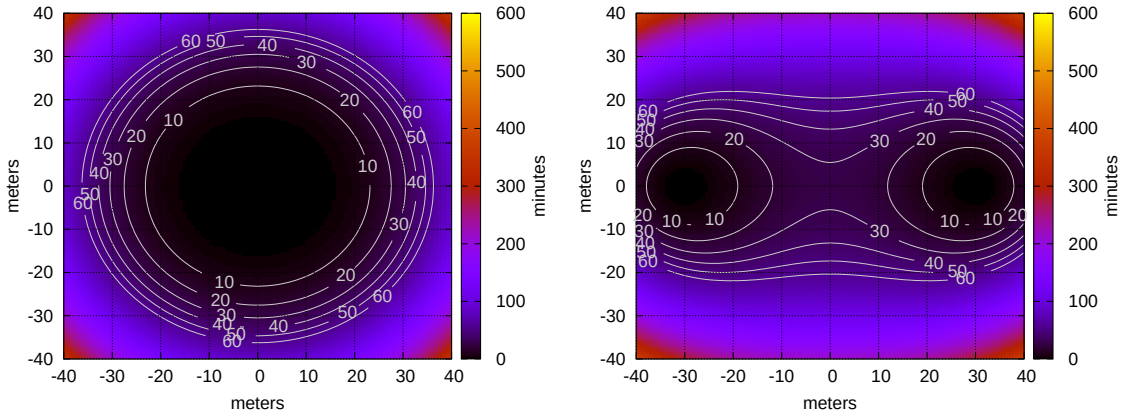

(a) **Monostatic Integration Time vs. Position (Center)**

(b) **Bistatic Integration Time vs. Position (Center)**

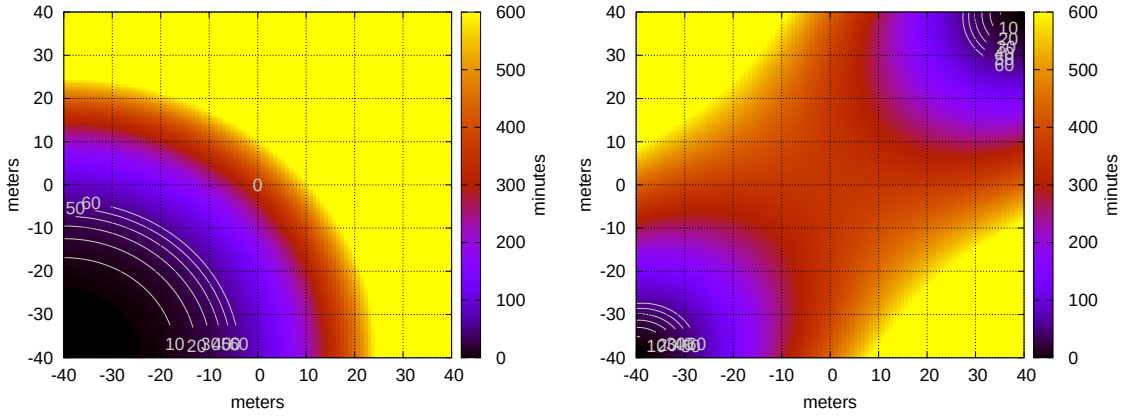

(c) **Monostatic Integration Time vs. Position (Corner)**

(d) **Bistatic Integration Time vs. Position (Corners)**

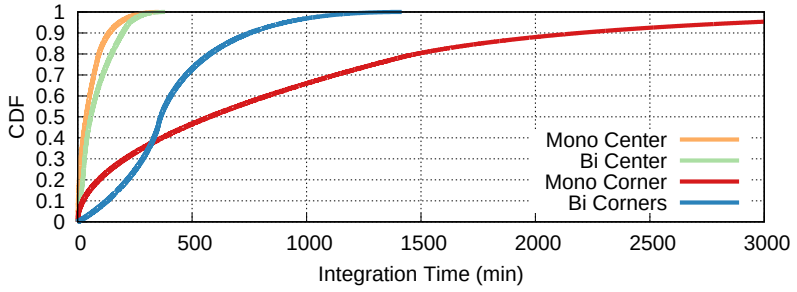

(e) **CDF of Integration Time**

Figure 5.6: **Integration Time (in Minutes) vs. Slocalization Tag Position for Two Different Anchor Arrangements.** The transmitting anchor can either be co-located (monostatic) or separated (bistatic) from the receiving anchor, each of which has its strengths and drawbacks. The monostatic arrangement suffers from high flash amplitude (the limited dynamic range of the RF frontend is overwhelmed by nearby high energy reflections) and inadequate spatial coverage in large areas. The bistatic arrangement results in a larger coverage area but requires time synchronization between the transmitting and receiving anchors, which are now physically separate.

the transmit and receive antennas, λ is the wavelength (in meters), and R_1 and R_2 are the distances (in meters) between the tag and the receive and transmit anchors, respectively. All gain and power figures are represented in decibels.

Recall that Slocalization's goal is the measure the distance between the tag and an anchor by determining the time of arrival of the reflection from the tag. A signal to noise ratio of approximately 26 dB is required for standard threshold-based leading edge detection techniques to accurately determine the tag's time of arrival. Using the example from Figure 5.5, with a transmitted signal power of -41.3 dBm/MHz and typical indoor settings of G_t,G_{bt},G_{br},G_r = 0 dBi, λ = 0.075 m, L_b = 1 dB, and R_1,R_2 = 5 m, the power received from the backscatter tag is -159 dBm/MHz.

5.4.1 Integrating Signal from Noise

In a stationary environment with no other signal sources, the ambient noise is approximately white and Gaussian, that is its integral over a long period of time is roughly zero. This observation leads to the *slow* in Slocalization: namely if you integrate a sufficient number of samples over time, you can extract the tag's signal from the channel impulse response. In section 5.7.5 of our evaluation we explore the impact of additional interference sources such as environmental motion or other ambient electronics, and show that these can be filtered out of the channel frequency response and do not significantly affect the performance of Slocalization.

To reliably recover reflection arrival time requires a 26 dB signal to noise ratio. Using the well-known interpretation of Johnson-Nyquist noise, we can express the noise as a function of integration time:

$$P_{dBm} = -174 + 10 \times \log_{10}\left(\frac{1}{t}\right) \tag{5.2}$$

where P_{dBm} is the noise power and t is the integration time. Continuing with the example in Figure 5.5 along with a receive noise figure of 10 dB, an integration time of less than one second is required to achieve a SNR threshold of 26 dB.

A variety of factors including obstructing materials, noise figure of the receiver, and nulls in the tag's antenna pattern can have a great effect on the parameters described in the backscatter path loss. Therefore, a significant margin of error must be applied in integration time to achieve high likelihood of tag detection in realistic indoor environments.

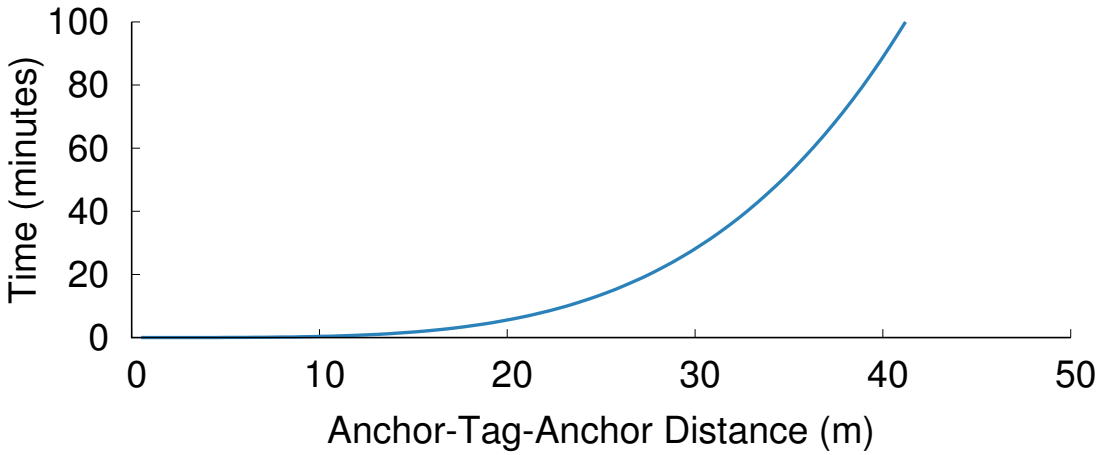

Figure 5.7: **The Time and Distance Tradeoff.** Finding a UWB backscatter signal requires integrating over multiple samples until the signal rises above the noise floor. The signal strength is directly correlated to the distance between the transmitter and tag. For a monostatic configuration (transmitter and receiver are the same anchor), this results in an exponential increase in required integration time with increasing distance.

5.4.2 Integration Time vs Distance

The distance between the tag and anchors can have a significant impact on the integration time required to recover the UWB signal. Figure 5.7 explores how integration time grows as the distance between the tag and anchor grows. For a monostatic configuration, where the transmitter and receiver are the same, this analysis finds that integration time grows exponentially and quickly becomes untenable for rooms larger than 40-50 m. A bistatic configuration, with separate transmitters and receivers, may perform better for some anchor placement and room scenarios, but the rough order of magnitude for room size (bounded by the anchor-tag-anchor path) will hold.

5.5 How to Measure Backscatter Transmissions: The Band-stitched UWB Receiver

The previous section described the UWB channel in theory. In this section, we explore the generation, manipulation, and recovery of UWB signals in practice.

Figure 5.8: **Slocalization Anchor Realization.** The Slocalization anchor uses a commercial off-the-shelf Ettus USRP N210 software-defined radio with the SBX 400-4400 MHz daughterboard. Set up as a bistatic measurement, each anchor transmits and receives at different frequencies. Through bandstitching, a highly accurate measurement of the UWB channel impulse response is derived, allowing for an accurate determination of the backscatter tag's CIR through time-domain correlation.

5.5.1 UWB Bandstitching

To address the limited availability of UWB hardware, Kempke et al. introduced the concept of a bandstitched UWB receiver [45]. The idea of bandstitching is that a more traditional and accessible narrowband receiver can capture a UWB sample by taking a series of narrowband samples at successive frequencies (3.33–3.36 GHz, 3.36–3.39 GHz. . .), add these samples together in the frequency domain, and then use this "stitched"-together sample to recover a high-fidelity UWB channel impulse response in the time domain.

5.5.1.1 Extending to Bandstitched Transmitters

For Slocalization we borrow Kempke's bandstitched UWB receiver architecture and extend the principle to UWB transmission. While this is modification is fairly straightforward, using bandstitching for both the transmitter and receiver introduces an additional system-level constraint that frequency hopping between the transmitter and receiver must be synchronized. This is trivial for the monostatic case, where the transmitter and receiver are the

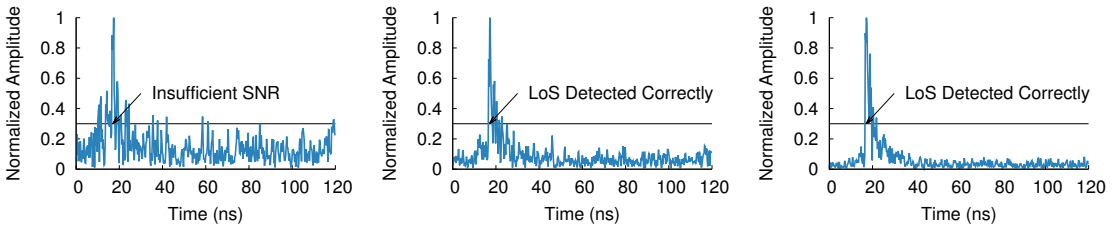

(a) 50 ms Integration (b) 250 ms Integration (c) 1250 ms Integration

Figure 5.9: **An example backscatter channel impulse response taken by Slocalization.** The 30% height of the CIR's leading edge is used to calculate the ToA of the line-of-sight path, necessitating sufficient SNR to resolve the leading edge. A number of integration lengths are shown for an example backscatter CIR. While 50 ms of integration time exhibits insufficient SNR to resolve the line-of-sight path, anything more than 250 ms shows sufficient SNR to resolve the backscatter CIR in this link scenario.

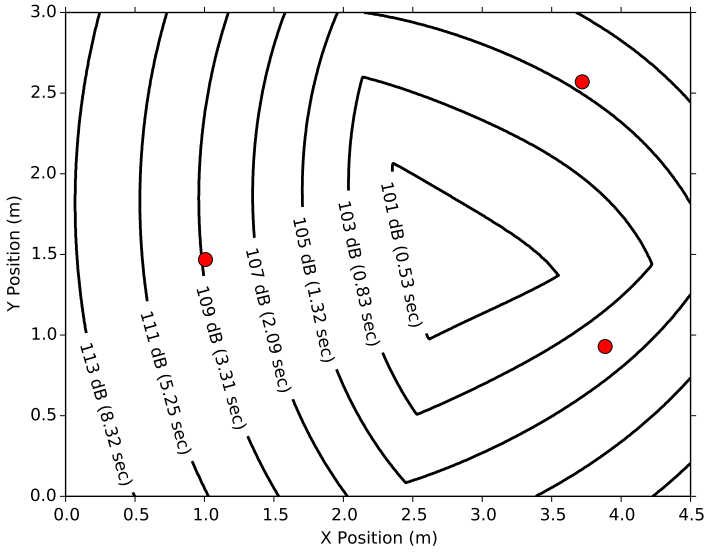

Figure 5.10: **Integration time versus tag position.** The minimum required integration time varies based on the maximum backscatter path loss observed between all recording anchors. A minimum signal to noise ratio of approximately 26 dB is required to accurately determine the tag's TDoA. Relative thermal noise decreases with increasing integration time, yielding the above dependence of integration time on tag position. Minimum integration time assumes 0 dBi gain at each antenna, 0 dB internal tag loss and 12 dB receiver noise figure. Longer integration times must be leveraged in instances where the tag or anchors fall within an antenna null.

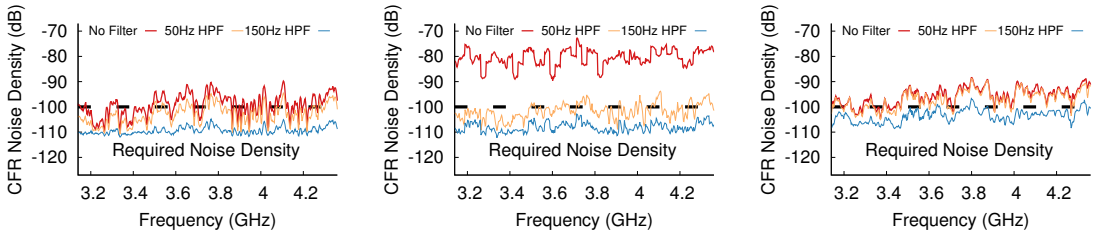

(a) CFR Noise Static (b) CFR Noise Walking (c) CFR Noise Fluorescents

Figure 5.11: **Effects on CFR from various dynamic environment processes.** In order to properly detect a backscatter tag, dynamic changes in the environment need to be compensated for. Here we can see the effects of different dynamic channel conditions on the CFR, the noise it imparts, and the effect of various filtering strategies, along with the required noise density requirement of a typical backscatter link with 100 dB of path loss. Walking around the environment imparts low-frequency noise which can be easily compensated through the use of a 50 Hz high-pass filter on CFR observations. Dynamic changes due to fluorescent lighting imparts higher frequency noise, requiring the use of a higher frequency high-pass filter to cancel. A control run shown in orange shows that even seemingly stationary environments observe CFR noise, likely due to noise internal to the software-defined radio. These effects require tag modulation to be high enough to not be drowned out by these common sources of noise.

same, but requires external synchronization for bistatic setups (distinct transmitters and receivers).

5.5.1.2 Comparison to FMCW

In WiTrack, Adib et al. face the same access to UWB problems as Kempke's Harmonium, however the WiTrack design requires a full transceiver. WiTrack's solution is to leverage the well-known frequency modulated carrier wave (FMCW) radar design to capture UWB range estimates [1]. The FMCW architecture has some drawbacks, however. It required WiTrack to design and implement a custom RF daughterboard for the USRP, which necessarily includes somewhat niche and expensive components, such as an ultra wideband VCO and a DDS.[1] In contrast, the Slocalization transceiver design runs on unmodified SDRs, closer emulating traditional narrowband radio architectures.

5.5.2 Backscatter Signal Recovery

To recover the tag's backscatter channel impulse response (CIR), anchors measure the backscatter channel *frequency* response (CFR) using standard narrowband CFR approaches.

[1] Voltage Controlled Oscillator; Direct Digital Synthesizer

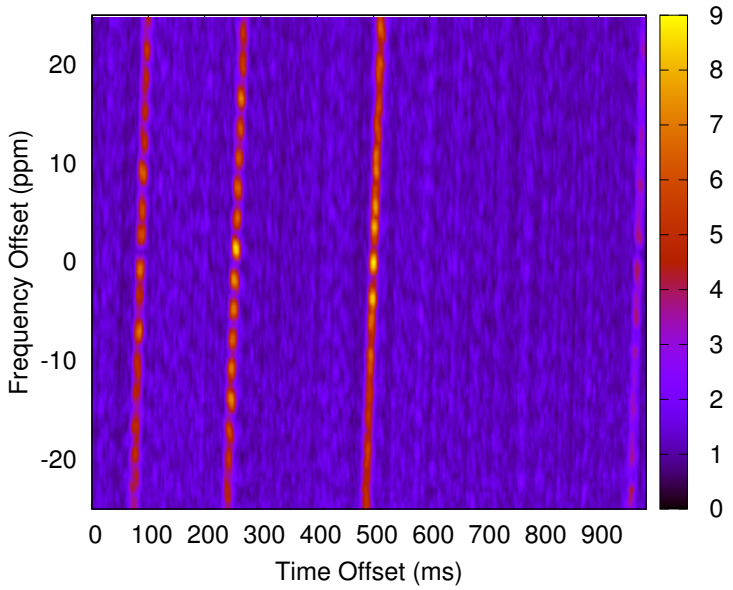

Figure 5.12: **Searching for Tags in Multi-User Settings.** To generate the backscatter tag CIR, the time offset and frequency offset of the backscatter modulation sequence must be determined. In the case of PN-coded backscatter transmissions, this search space can be quite large. This shows the resulting correlation search space for a PN code of length 63 transmitted with a period 983 ms. Three tags can be observed after an exhaustive search is performed. The peak values for each tag are used to accurately correlate and reconstruct their corresponding backscatter CIRs.

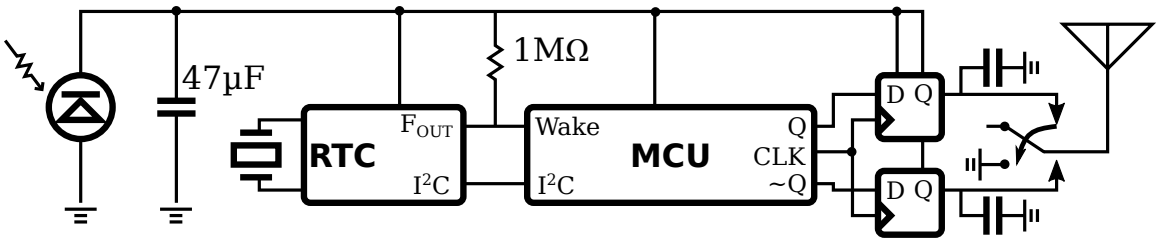

Part	MPN	Quantity	Cost (USD qty 1k)
MCU	STM32L051K8T6	1	1.80
Antenna	AH086M555003-T	1	1.57
Solar Cell	AM-1417	1	1.44
RF Switch	UP62422TK	1	0.71
RTC	AM0805AQ	1	0.55
Crystal	ABS07-32.768KHZ-7-T	1	0.38
Flip Flop	74LVC1617S	2	0.09
Passives			0.16
PCB		1	1.00
Total			7.70

Figure 5.13: **Slocalization UWB Backscatter Tag Schematic.** To maximize flexibility, we utilize a low-power MCU in place of a shift register for our tag implementation. To achieve the requisite oscillator stability at minimal power budget, we use an ultra low power real time clock from Ambiq. To maintain a minimal tag power draw, we sleep the MCU between (potential) bit flips, requiring a pair of flip flops to drive the RF switch. We include a simple energy harvesting frontend made up of a 47 µF energy storage capacitor and an indoor photovoltaic cell.

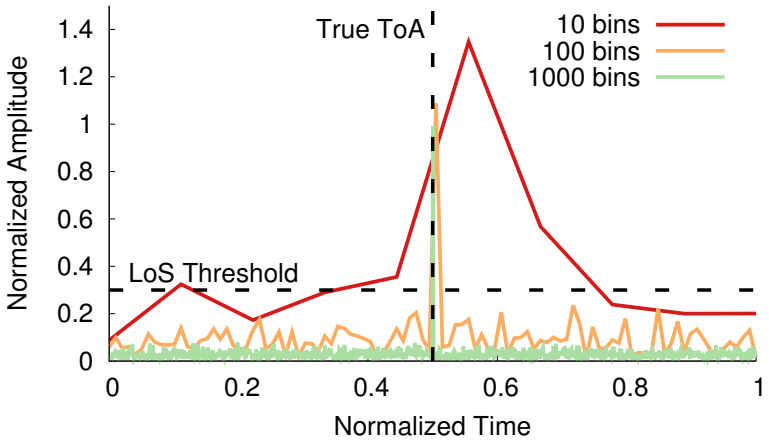

(a) CIR vs. Number of Bandstitched Bins

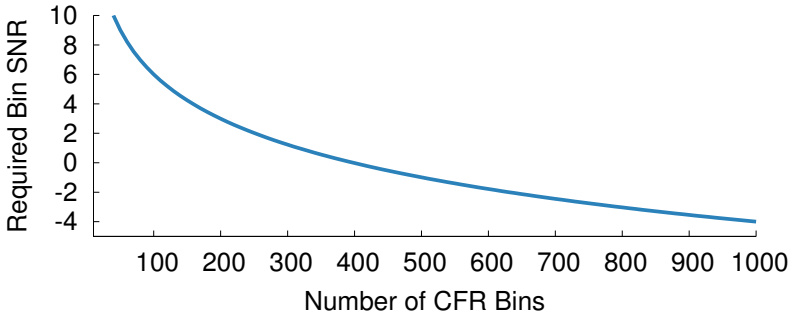

(b) Required CFR SNR vs. Number of Bandstitched Bins

Figure 5.14: **Processing Impacts Precision.** The introduction of more bandstitching bins not only contributes to better CIR resolution from greater utilized bandwidth, but also improves the CIR SNR, given the same integration time for each CFR bin. The increase in SNR is due to the coherent contribution of many, noisy CFR bins. In the case of the single-path case, the CIR SNR increases by $10 \times log_{10}(n_{bins})$. A CIR SNR of 26 dB is required in order to reliably observe the LoS path.

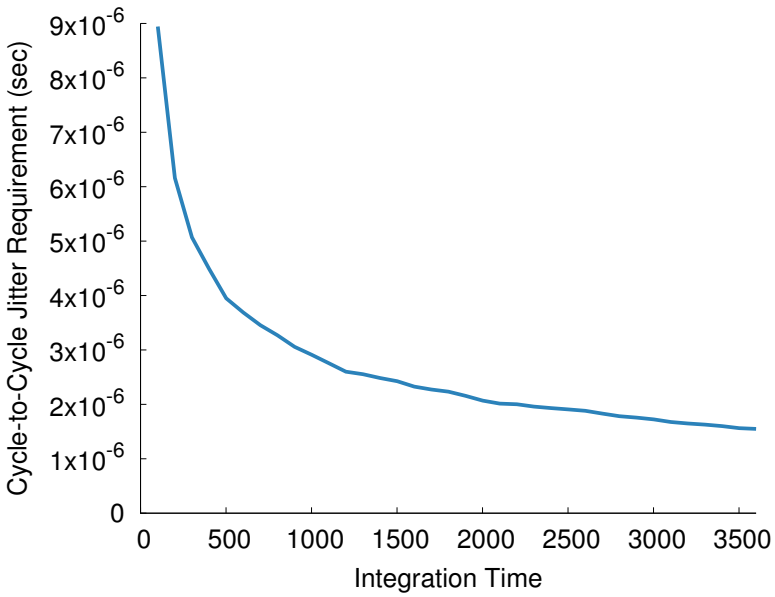

Figure 5.15: **Crystal Stability Limits Distance.** To recover tag signals, the UWB receiver must be able to correlate the tag pulse train. This requires the pulse generation to remain stable over the course of integration window at the receiver. As the required receiver integration time increases, the allowable tag jitter falls.

To differentiate the backscatter CFR from the direct CFR, successive measurements of the CFR are correlated to the tag's transmit sequence.

To successfully extract the tag's signal, the tag's transmit sequence must follow two characteristics. First, the sequence must have a zero mean (an equal number of zeros and ones). This ensures that no portion of the direct CIR is present after correlation. Second, the sequence must employ a modulation rate higher than that of other dynamic sources within the environment. To meet these requirements, Slocalization mixes the transmit sequence with a pattern of the form $sgn\left(sin\left(\omega * t\right)\right)$.

The receiver must correlate against both tag frequency and sequence offset to recover the tag's CIR. A running count of time is kept between observed bandstitching steps to account for tag drift while switching bands. A complete estimate of the tag's CIR can be computed after an entire UWB sweep. To achieve sufficient backscatter SNR for a single backscatter CIR, careful consideration of the dwell time at each band is required.

Figure 5.14 describes the tradeoff between the number of CFR bins and the CIR variance. Due to the coherent summation of CFR bins the required SNR for each CFR bin decreases with an increasing number of bins. The coherent summation of N bins yields a $10 \times log_{10}\left(N\right)$ increase in CIR SNR. To achieve an approximate 26 dB CIR SNR[2], a CFR

[2]26 dB of CIR SNR yields a negligible number of false positives in CIR ToA detection

bin SNR of $26 - 10 \times log_{10}(N)$ is required. This relates directly to the required dwell time at each band.

The number of bandstitching steps along with the dwell time at each step defines the time to complete a full UWB sweep. Various methods can be employed to increase the UWB sweep rate. Either the instantaneous bandwidth can be increased through the use of higher sampling rate ADCs or multiple bands can be observed simultaneously through observation across multiple center frequencies. The current Slocalization design employs 25 MHz of instantaneous bandwidth utilizing one RF receive frontend, yielding 49 steps with two seconds of integration time at each step to generate 1.225 GHz of UWB sweep bandwidth over the course of 98 seconds of total integration time.

With 2 seconds of integration time at each step, integrated thermal noise is -177 dBm per Section 5.4.1. A reasonable receive noise figure of 10 dB increases the effective noise floor to -167 dBm. The 490 total CFR bins measured across the entire UWB bandwidth yields a processing gain of $10 \times \log_{10}(490) = 26.9$ dB. Therefore, to achieve 26 dB of CIR SNR, a CFR SNR of -0.9 dB is required. Assuming a UWB-compliant signal with -41.3 dBm per bin, this gives an allowable total backscatter path loss of $(-41.3) - (-0.9) - (-167) = 126.6$ dB.

In order to support long integration times at the anchor, timing jitter in the tag's modulation sequence must be kept to a minimum. Over the course of the anchor's integration period, the modulation sequence must not deviate by more than 1/4 bit period from the average rate. Figure 5.15 shows the allowable signal jitter vs. integration time for the 256 Hz tag modulation rate evaluated in this paper, derived through Monte Carlo simulation. RC jitter for the AM0805 (a representative frequency source) is approximately 500 ppm, necessitating the use of the higher-power crystal mode to maintain code coherence.

5.6 Slocalization Implementation

As a backscatter system, realizing Slocalization does not require many components. However, due the sensitivity of the backscatter channel and a focus on minimal power draw, careful selection of components is required to maximize the potential of Slocalization.

5.6.1 Tag Design

The ultra-low power requirements for Slocalization's UWB backscatter tag allows for a batteryless design, freeing itself from the periodic recharging or replacement maintenance typically associated with battery-based designs. The following sections detail the imple-

mentation of the Slocalization tag's backscatter and energy harvesting circuitry.

5.6.1.1 Backscatter Frontend

The principle of backscatter relies on the ability to modulate an antenna's impedance between varying states. To achieve this, the Slocalization tag leverages a SPDT RF switch. One input to the RF switch is tied to ground (through an AC coupling capacitor) while the other is left floating, allowing for a 180 degree phase shift between two states. The output of the RF switch is connected to the tag's UWB antenna.

The Slocalization tag uses the UPG2422TK RF switch due to its desirable characteristics of a) minimal insertion loss, b) low power operation, and c) low switching voltage. The two RF switch control lines are driven by D flip flops which are latched using a low-power MCU. The MCU modulates the tag's PN backscatter sequence and is used to emulate the functionality of a simpler shift register-based design while also facilitating greater experimental flexibility.

For a stable frequency source, the Slocalization tag leverages the AM0805, an ultra-low power RTC that draws only 50 nA to provide a stable frequency reference. This RTC is configured to provide the MCU with the appropriate backscatter modulation rate (256 Hz in our implementation).

5.6.1.2 Energy-Harvesting Details

For the purposes of this evaluation, the Slocalization tag achieves light-based energy harvesting through the use of an attached 35 mm $\times$ 14 mm photovoltaic cell. With an open-circuit voltage of 2.4 V in indoor conditions, it provides sufficient voltage to run all embedded Slocalization tag electronics. An energy storage capacitor is used to provide sufficient high-current switching demands from the embedded MCU.

Any number of alternative energy harvesting designs have been shown to provide power delivery capabilities above the 1 µW necessary to power the Slocalization tag. Of particular interest are that of narrowband RFID techniques, which may be able to leverage the UWB antenna already included on the tag.

While the Slocalization tag has been shown to achieve UWB backscatter on an extremely constrained power budget, we believe there is significant headway achievable with future designs. For instance, the functionality demonstrated by the MCU can be replaced with a linear feedback shift register. Additionally, the D flip flops and RF switch all have the potential to be integrated onto a single die. Future work will allow for a deeper analysis of the size and power achievable through a custom silicon design.

5.6.2 Anchor / Network Design

Deployed Slocalization infrastructure consists of a number of anchors, such as that shown in Figure 5.8, each of which are synchronized with a shared clock and connected to a host computer using gigabit Ethernet. The host computer orchestrates the coordinated reception of a single bandstitching anchor's transmissions across the entire UWB bandwidth.

Transmit data is fed to the designated TX anchor as a repeating sequence of twenty IQ samples, chosen as a sequence which minimizes dynamic range and maintains equal amplitude across the 25 MHz of occupied bandwidth. Due to the repetitive nature of the signal, this sequence is designed to generate twenty CFR peaks across 25 MHz, calibrated to a transmit amplitude abiding by the FCC requirement of -41.3 dBm/MHz.

The rest of the anchors involved in the network feed received IQ samples back to the host PC for post-processing. An initial real-time integration step averages out high frequency effects. For example, at 20 samples per repetition and 25 Ms/s, a CFR update rate of 1.25 MHz is achievable, but not useful for the low tag modulation rates utilized by Slocalization. This initial integration step repeatedly integrates the 20-sample sequence *one thousand times* before dumping the averaged IQ data. This 1000× decimation yields a CFR update rate of 1.25 kHz, enough to cover the modulation rates utilized by Slocalization while minimizing additional signal processing complexity.

Once the UWB sweep is performed once, the role of all participating anchors is rotated to calculate each TDoA estimate in turn. Future work will explore the simultaneous transmission and reception of multiple anchors to minimize required end-to-end integration time.

5.6.3 Processing

All signal processing for Slocalization is performed in MATLAB using raw USRP data. The following sections walk through the processing step by step.

5.6.3.1 Data Parsing and Trimming

Averaged IQ data as described in section 5.6.2 is stored in a format which includes tagged metadata identifying the precise time and target of retune events occurring within data logging operations. These retune events are used to segment the IQ data into separate bandstitching snapshots. After IQ data segmentation, the first 80 ms of each step are then trimmed to allow the receiver's RF PLL to settle to the newly-tuned frequency.

5.6.3.2 Clock Ambiguity Resolution

When distributing time via a shared reference, a number of PLLs and their effects must be considered. First, time is distributed as a 10 MHz signal to each anchor. From there, it is multiplied by $10\times$ to provide clocking internal to the USRP. This reference is then divided by $4\times$ to provide the reference for the transmit/receive RF PLL. Depending on the random timing introduced through the power-on sequencing internal to each radio, the phase of the final 25 MHz signal can be offset in time between anchors. A signal processing step in software measures the phase difference incurred between received bands and corrects for any phase offset incurred.

5.6.3.3 CFR Calibration

Each adjusted IQ sequence is then compared to the expected receive sequence through a calibration step. Direct-connection calibration recordings between pairwise anchors are performed to accurately compensate for any phase offset incurred during RF signal generation and reconstruction. CFR measurements are directly generated through deconvolution with the calibration dataset.

5.6.3.4 Direct CIR Acquisition

Direct CIRs between anchors are generated through assembling the full UWB CFR. These direct CIRs are later used to calculate relative TDoA to the backscatter tag.

5.6.3.5 Backscatter CIR Acquisition

The backscatter CIR is calculated by correlating the CFR with the tag's known backscatter modulation sequence. The search consists of a 2D search including tag modulation frequency and sequence time offset.

5.6.3.6 TDoA Estimation

Once the direct and backscatter CIRs are known, the TDoA can be estimated by calculating the difference in ToA between the direct and backscatter CIRs. Our implementation uses the 30% height of the tallest peak to estimate the time-of-arrival of each CIR, a fairly standard approach.

5.6.3.7 Ellipsoid Localization

Multiple TDoAs to a single tag are leveraged to determine the tag's location. The TDoA estimates along with the known positions of participating anchors define ellipsoids of possible tag positions. The minimum mean squared error is determined through a gradient descent algorithm to determine the best-fitting position estimate given the measured TDoAs to each reporting anchor pair.

5.7 Slocalization Performance Assessment

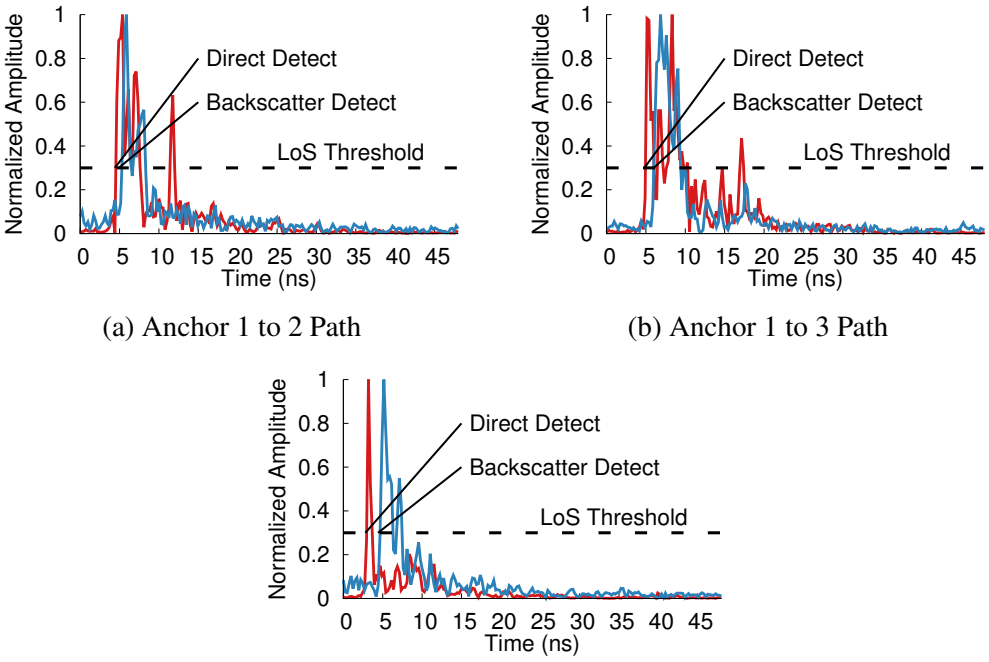

(a) Anchor 1 to 2 Path

(b) Anchor 1 to 3 Path

(c) Anchor 2 to 3 Path

Figure 5.16: **Direct- and Backscatter-Path Channel Impulse Response Measurements.** Each CIR is calculated using 1.225 GHz of bandstitched ultra-wideband measurements. For each pairing of anchors situated throughout the environment, the channel impulse response is calculated. Multiple time difference-of-arrival backscatter measurements are necessary to determine a tag's 3D location accurately.

In this evaluation we aim to prove the viability of Slocalization. We first demonstrate the system's capability to recover TDoA estimates from a backscatter signal. We then demonstrate end-to-end localization of a Slocalization tag in a complex indoor environment, finding it achieves 30 cm average error across an array of points. We then evaluate some of the underlying Slocalization components and answer the question of how Slocal-

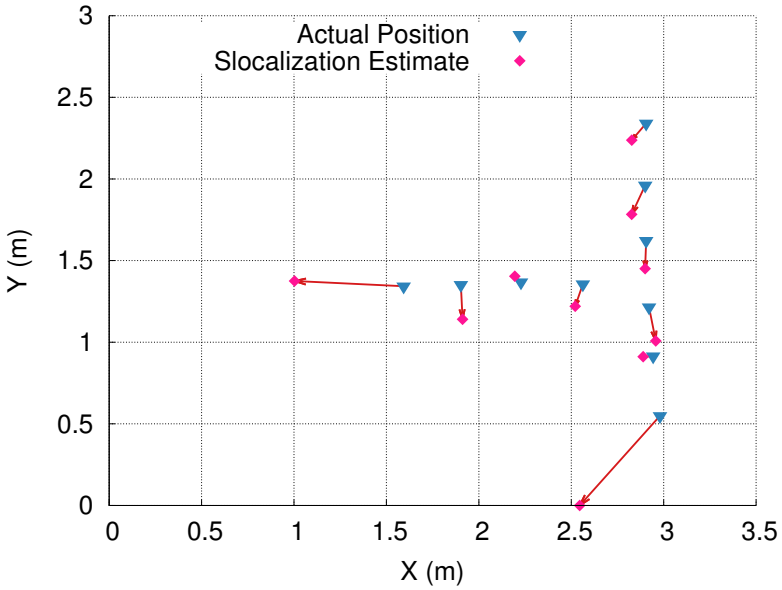

	Min	Max	Mean	Median
2D XY Euclidean Error	0.05	0.70	0.25	0.18
3D Euclidean Error	0.08	0.70	0.30	0.26

Figure 5.17: **Slocalization Performance Evaluation.** Ground truth vs. estimated tag position in a 4.5 m $\times$ 3 m $\times$ 2.3 m interior room. A number of fixed locations are chosen for the Slocalization tag, and the difference between the calculated position and the true position are shown. Slocalization is able to achieve 30 cm of average 3D error using sub-microwatt tags across the entire evaluation space using only 98 seconds of integration time at each location.

ization can handle and reject environmental interference. Finally, we show that we can distinguish and recover ranging information from multiple Slocalization tags transmitting in parallel in the same environment.

5.7.1 Can Slocalization Measure TDoA?

We begin by investigating the viability of Slocalization. We set up three anchors configured for bistatic ranging (separate transmitter and receiver) and a single Slocalization tag. In Figure 5.16, we show the recovered channel impulse response for the Anchor $1 \rightarrow 2$, Anchor $1 \rightarrow 3$, and Anchor $2 \rightarrow 3$ paths. The Slocalization system can clearly identify peaks for both the direct and backscattered path for all anchor pairings. This time difference of arrival (TDoA) coupled with known 3D positions of anchors can be used to localize the Slocalization tag.

5.7.2 3D Location Estimation

We next investigate the quality of the location estimates provided by Slocalization. We set up Slocalization in a 4.5 m $\times$ 3 m $\times$ 2.3 m interior room and place the tag in 10 different locations throughout the room. We configure the bandstitching sweep to dwell for 2 s at each of the 49 measured bands, requiring 98 s total for each location fix, an update rate of approximately 10 mHz. Figure 5.17 shows the results of this experiment and finds that the Slocalization system is able to achieve an average error of only 30 cm across all 10 locations.

5.7.3 Integration Time

The integration time is the key factor that determines how slow Slocalization runs. Because the signal received from the tag is well below the noise floor, the Slocalization system needs to integrate numerous samples of the environment over time to extract the tag's signal. Recall, the goal is to be able to accurately detect the leading edge of the pulse reflected by the tag, as the time offset of this edge yields the distance between the tag and anchors. Figure 5.9 looks at the effect of varying this integration time for a sample link.

In this experiment the tag-anchor-tag distance is just shy of 5 m, which allows an integration time as short as 250 ms to successfully recover the line of sight path. Note that 250 ms is only the integration time for one slice of the ultra wideband spectrum. Bandstitching requires 250 ms of dwell time at each of the 49 frequency slices, thus requiring 12.25 s to completely resolve position.

5.7.4 Integration Time vs Position

Figure 5.7 presented a theoretical analysis of the impact of distance on integration time. In Figure 5.10, we evaluate the actual integration time across our evaluation space. The required integration time for a given position is dictated by the worst case (most path loss) from any tag-anchor-tag configuration. In practice, this means that anchor placement can have a significant effect on the best, average, and worst case localization time for positions throughout a room.

5.7.5 Environmental Noise

A principle design goal of Slocalization is the ability to localize a static tag. This design goal extends to the ability to localize the tag in a static environment. However, in many real-world scenarios, while the localization target may be stationary, the environment often is not. A non-stationary environment will appear as noise in the channel frequency response (CFR). In Figure 5.11a we capture the CFR in a static environment and then add various perturbations and compensation strategies.

We first consider the obvious environmental noise source for indoor spaces, namely people moving throughout the environment. In practice human beings do not move quickly in physical space, and Figure 5.11b shows that the simple addition of a 50 Hz high-pass filter is able to remove most of the CFR noise created by people moving about the space.

The next source of noise Slocalization must deal with is that emitted by ambient devices in the space. In Figure 5.11c we find that the fluorescent lighting in our office building emits significant noise not successfully filtered by the 50 Hz filter added for removing human motion. Raising this filter to 150 Hz successfully removes the noise introduced by the lighting, facilitating Slocalization.

It is in Slocalization's interest to keep this filter value as low as possible. The primary energy cost for the tag is turning on and off the antenna load switch, thus the lower the switching frequency, the lower the tag's active power draw. In practice we have not found other significant interference sources above 150 Hz testing in both a traditional office setting and a home environment. We set the tag oscillation frequency to 256 Hz to balance active power draw and detectability.

5.7.6 Multiple Tags

The Slocalization design includes PN codes to allow the anchor infrastructure to distinguish multiple tags. Figure 5.12 places concurrently transmitting Slocalization tags in the

environment. The Slocalization system is able to cleanly distinguish each tag and localize it independently of the others.

5.8 Summary

In this chapter, we have described the design and implementation of a system which is able to accurately track the position of stationary, energy-harvesting tags. Such a localization architecture has been shown to be applicable to the asset tracking domain due to the minimal maintenance requirements and minimal power generation capability of energy harvesting technologies.

The Slocalization tag consumes less than one micro-watt of power and consists of an UWB antenna, RF switch, and RTC, along with a microcontroller used only once to configure the RTC's settings. Power is applied from a solar cell, but can also be supplied from any number of other identified energy harvesting technologies.

The Slocalization anchors are realized using commodity radios configured to generate ultra-wideband channel impulse response measurements. Due to self-interference issues, the anchors measure the tag's backscatter CIR between anchors by measuring the minute effects of the tag's backscatter modulator on the channel impulse response observed between anchors.

This section concludes our exploration of the applications of personal navigation, robotic navigation, and asset tracking capabilities. The next section will focus on the troubles which will arise when deploying the infrastructure necessary to perform these functions along with potential solutions to these issues as well.

CHAPTER 6

Enabling Broader Applicability: Global Time Synchronization and Scheduling

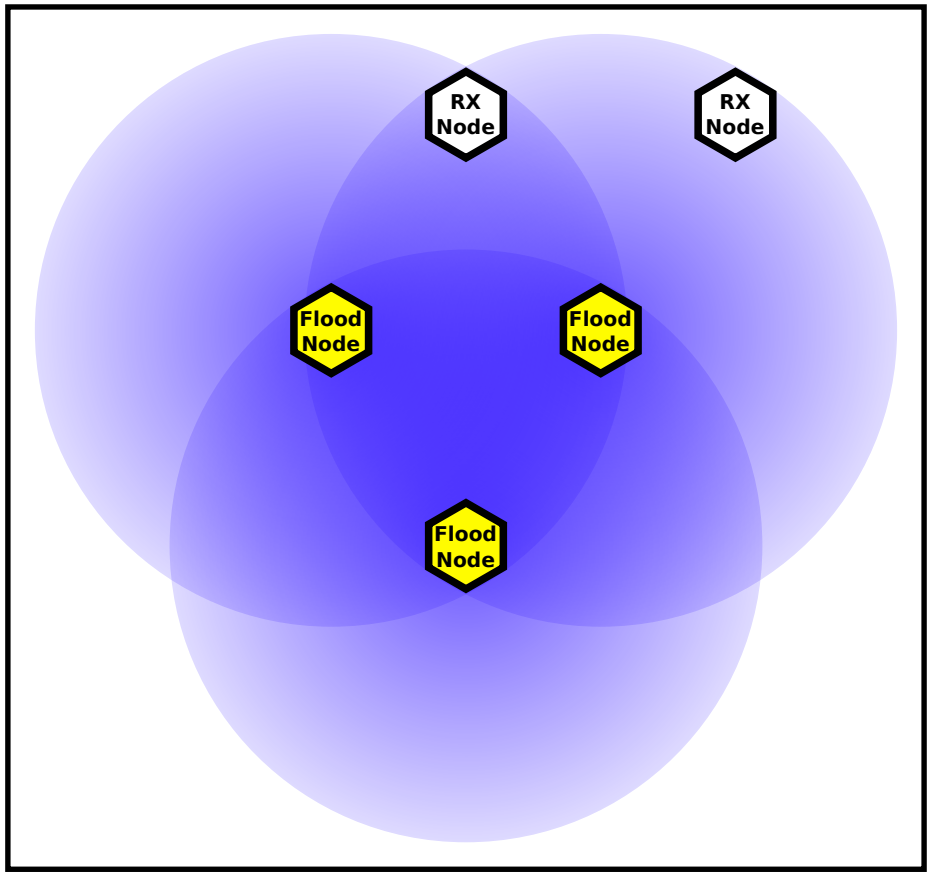

Figure 6.1: **SurePoint CIR Flooding Concepts.** If sufficient clock synchronization is attainable, the simultaneous transmission of multiple transceivers mimics the effects of multipath propagation. Furthermore, the combined CIR contains more power, giving rise to more reliable message communication.

Up until now, no method for orchestrating the localization of multiple tags in a shared

environment has been described. SurePoint tags request ranges to nearby anchors at will, and this will inevitably result in collisions in highly contested environments.

For this reason, further functionality has been added to the SurePoint protocol to enable the global time synchronization and scheduling of all tags and anchors participating in the network. This allows for a graceful degradation of performance with the addition of tags operating in the same area, as there will be no chance of network collisions during localization operations.

This functionality is enabled by the application of constructive interference, which allows for the simultaneous transmission of many tags at the same time to aid in increasing link reliability and enabling the adoption of physical-layer protocols such as LWB in the ultra-wideband space. The rest of this chapter will detail these concepts of constructive interference when applied to UWB and how it can be used to provide accurate global time synchronization and scheduling.

6.1 Constructive Interference With UWB

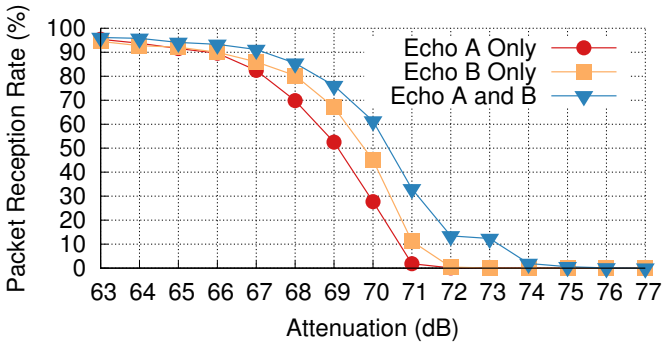

Figure 6.2: **Constructive Interference With UWB.** To evaluate whether UWB achieves constructive interference we place three nodes, an initiator and two echoers, into Faraday cages and connect them with variable attenuators. When two nodes reply concurrently, the packet reception rate rises, validating that constructive interference is viable in the 802.15.4a UWB channel.

Constructive interference – or the simultaneous transmission of multiple nodes to achieve better link reliability – has been demonstrated and leveraged in the narrowband case [22] by protocols such as the Low-Power Wireless Bus (LWB) [26] which is built on top of the 2.4 GHz 802.15.4 radio and Glossy floods [27]. In the case of narrowband transmissions, the combined transmission of many nodes has been utilized to provide time synchronization primitives in networks of three or more nodes. In the case of narrowband transmissions,

packets must be kept short in order to avoid the enveloping effect which occurs when multiple sinusoids are summed together with varying frequencies (due to either crystal variances or time synchronization mismatch).

The situation is slightly different with ultra-wideband transmissions. Due to the sparse nature of the ultra-wideband channel, transmissions which are separated by only $1/BW$ (down to two nodes) constructively interfere *perfectly* (3 dB) in a single-path environment. Reduced performance may be seen in heavy multipath environments where the delay spread is large, but will perform *at least as well* as in the narrowband case.

6.1.1 Constructive Interference Results in UWB CIR Superposition

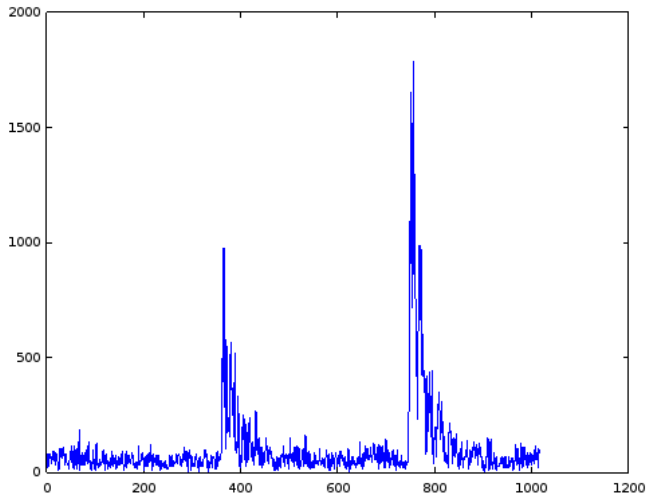

Figure 6.3: **Measured CIR Superposition.** The effect on the CIR for two simultaneous UWB transmissions spaced apart by 400 ns for clarity. Total signal power is increased by approximately 3 dB, aiding in the constructive interference (and thus further reach) of these closely-synchronized transmissions.

Figure 6.3 shows the effect on the derived CIR of two nodes simultaneously transmitting (with an artificial offset of 400 ns) in a typical indoor setting. The CIR now contains the sum of the two channels observed between the receiver and two transmitters. The overall power contained within the CIR is equal to the sum of the power in each separate CIR. In future work, multiple time-of-arrival estimates could be extracted from the combined CIR assuming the delay spread of the two CIRs is small enough to ensure minimal overlap.

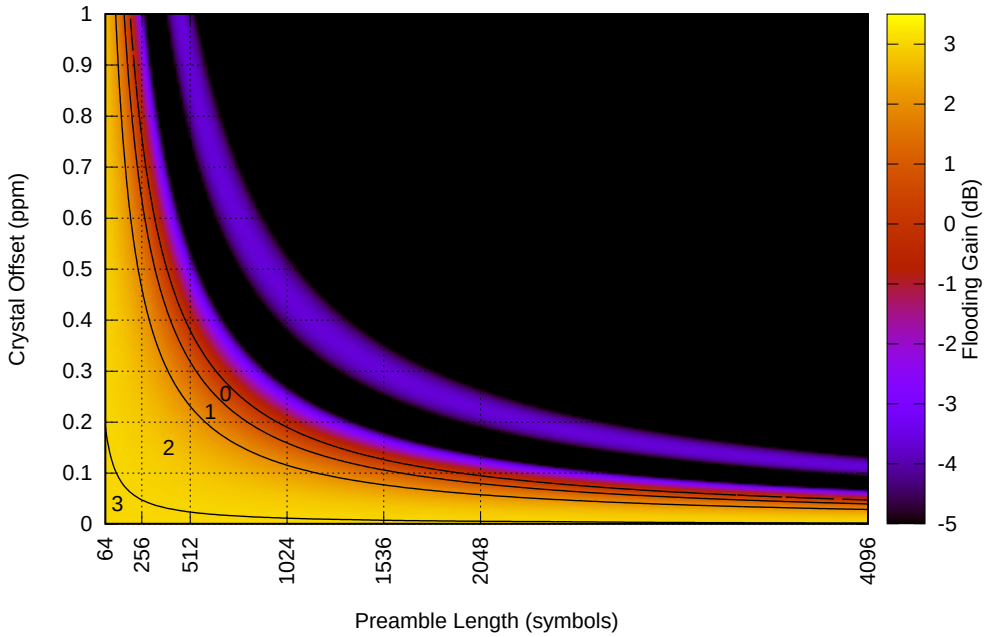

Figure 6.4: **Flooding Efficacy vs. Crystal Synchronization.** Channel impulse response power vs. carrier offset for two equidistant nodes transmitting simultaneously in a single path environment. Modest crystal inaccuracies can result in a significant degradation of the positive effects of ultra-wideband flooding which include increased communication range and robust data transmission.

6.1.2 Higher Crystal Accuracy Requirements for UWB CI

In order for the two received CIRs to constructively interfere, the carrier frequency of each must be close enough to the receiver local oscillator so that the difference in phase between the two between the start and end of preamble accumulation is less than π radians. Figure 6.4 shows the relative signal power improvement vs. carrier offset for the simultaneous transmission of two flooding UWB transmitters.

The maximum preamble time specified by the 802.15.4a specification is approximately 4 milliseconds. This correlates to an approximate crystal accuracy requirement of 0.05 ppm to ensure adequate constructive interference capabilities. This accuracy requirement is much more stringent than normally obtainable through inexpensive crystal oscillators, so a different method of achieving this synchronization accuracy between nodes is required.

6.2 Extending Constructive Interference: UWB Flooding

Constructive interference has been found to be helpful in providing a simple means to providing a flooding primitive in prior work on narrowband wireless networks. The Low-

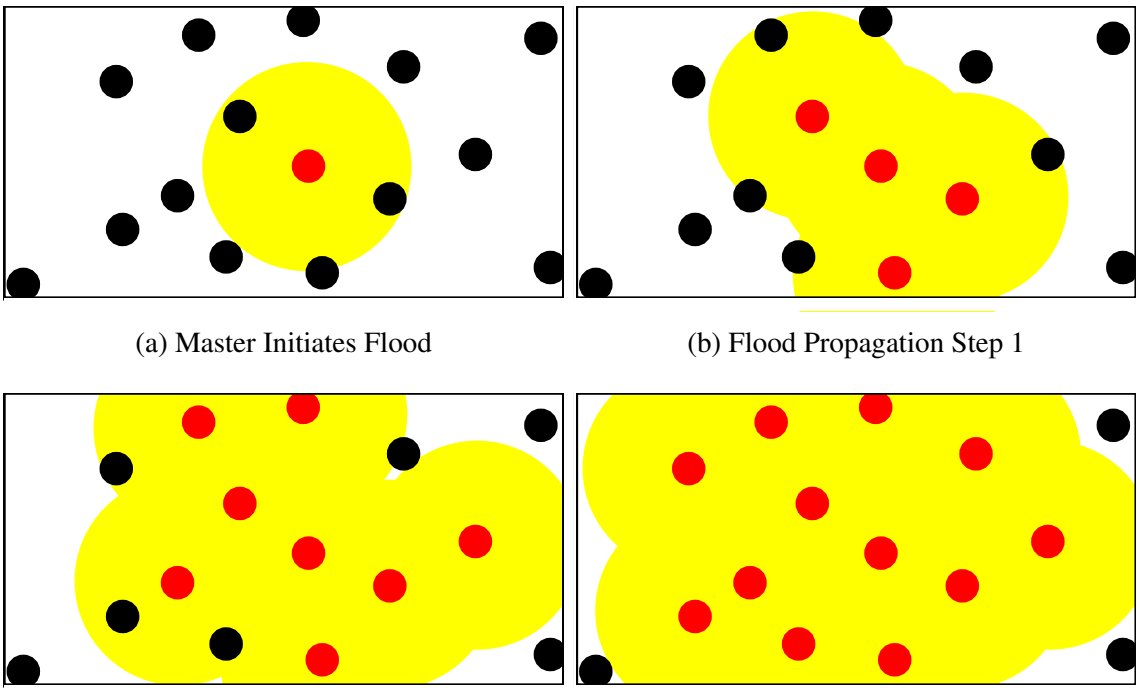

(a) Master Initiates Flood

(b) Flood Propagation Step 1

(c) Flood Propagation Step 2

(d) Flood Propagation Step 3

Figure 6.5: **Time Succession of a UWB Flood.** The UWB flood is kicked off by request of a master node. Each time step thereafter, any node which has received any of the previous flood transmissions perpetuates the flood until an incrementing depth counter is reached, corresponding to the maximum expected depth of the UWB network.

Power Wireless Bus (LWB) utilizes constructive interference to extend wireless reach. Floods are initiated by a single node, then all nodes which have received the flood initiation message then re-transmit the same message at a fixed time offset. This is perpetuated for a number of successive re-transmissions equal to the maximum expected depth of the wireless network.

6.3 Problem: Uncoordinated Tag Transmissions

SurePoint was developed to create a hardware and software abstraction for localization operations along with exploring the ways in which diversity can be leveraged to gain an improvement in localization accuracy. An extension of these capabilities is required to realize applications requiring the use of many, transient tags.

In prior work, the Low-Power Wireless bus is used to provide a global scheduling scheme with all communications based on narrowband floods. An extension of the Low-Power Wireless bus to the Ultra-Wideband domain is a simple extension which requires relatively little effort. The following section describes the details of the specific implementation for the global time synchronization and scheduling of many SurePoint tags.

6.4 Solution: Provide a Global Time Synchronization Methodology Based On UWB Floods

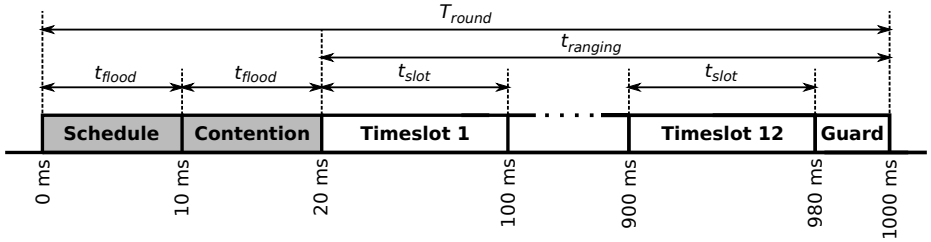

Figure 6.6: **SurePoint Multi-Tag Protocol.** Every round begins with a scheduling flood sent out by the master anchor. This schedule includes the round duration, contention slots for joining the network, and the schedule for each tag. To improve update rate, at the cost of network join latency, the round duration is much larger than the duration of a SurePoint range event, which means the same tag may be scheduled multiple times in a single round. In this example, the round duration is fixed at 1 Hz with 80 ms timeslots allowing up to twelve ranging events per second to one tag or one ranging event per second for twelve tags.

All scheduling and time synchronization operations are orchestrated by a single, 'master' node whose role is dictated during network setup. This master node sends a periodic synchronization flood which allows all nodes participating in the network to reference a common time base. The remaining timing assignments between scheduling floods are shown in Figure 6.6.

Along with providing global synchronization through the schedule flood, a snapshot of the current schedule and any newly-scheduled tags are included in the schedule flood data. After the schedule flood, a time slot is reserved for a contention period. This contention period allows for unscheduled tags to request a time slot during the schedule. Any nodes which are currently unscheduled pick a random time offset within the contention slot to initiate a flood. The first flood received by the master node during the contention slot is then scheduled and notified during the scheduling flood in the next schedule round.

All remaining time slots are reserved for ranging transmissions to and from the node which is assigned during that time slot. At most one tag is allowed to perform localization operations within each assigned time slot in order to avoid collisions of packets used during two-way ranging operations.

6.4.1 Joining the Network

Once SurePoint has achieved accurate network-wide time synchronization through the use of periodic Glossy floods, LWB defines a contention timeslot for use in requesting schedule assignments. In the case of SurePoint, a schedule request is sent by each unscheduled tag at a random time offset within this slot. Those who are listening during this time interval (anchors and scheduled tags) proceed by perpetuating the flood of the first schedule request they see. This flood is sent with the intention of it reaching the master anchor who will add it to the schedule. Note that from the perspective of a new tag, there is no difference between simply sending its packet after a different flood has started or two tags sending simultaneously and spawning competing floods. The only important metric is whether the master anchor hears a tag's request.

New schedule assignments are sent in the synchronization flood that begins the next round. To keep synchronization packets short in the common case, each tag is assigned a unique id. When first scheduled, the synchronization flood packet sends out a mapping of tag EUI to id. At most one new scheduling assignment can be made per round in the interest of keeping the packet length short (recall that longer packets require better crystal calibration). The schedule itself is defined by the number of tags (schedule length), the optional assignment of a new tag EUI to a localization timeslot, and a bit array containing

the valid tags.

If the schedule request flood is not received by the master or a schedule assignment is not heard by the requesting tag, the tag will retry each round until a valid schedule assignment is received.

6.4.2 Steady State Operation

Once all tags within reach of the SurePoint network are scheduled, the network enters a period of steady state operation where each tag performs localization tasks within one or more fixed timeslots after the contention period. After the contention slot, the remaining time in T_{round} is allocated to $t_{ranging}$, which is divided into $t_{ranging}/t_{slot}$ slots. Each tag counts round timeslots modulo the total length sent in the schedule packet. Whenever the round timeslot matches the tag's assigned timeslot, the tag starts the SurePoint ranging protocol to nearby anchors. If there are fewer tags than time slots, some tags will have the opportunity to range multiple times per round.

6.4.3 Leaving the Network

Due to the inherent mobility of nodes participating in localization operations, nodes will frequently go out of range of the SurePoint network or may wish to stop ongoing localization operations altogether if knowledge of their position is no longer required. For this reason, SurePoint implements two methods for removal of nodes from the network. Those nodes which would like to give up their timeslot can issue a disconnection request during the contention round through the use of a dedicated flood back to the master. Each node's timeslot is also given a lifetime, and the master will de-schedule those nodes which haven't been heard from during that period. In order to avoid disconnection, nodes wishing to continue on-going localization operations can re-issue their schedule request prior to the lifetime limit through use of a repeated scheduling flood back to the master node.

6.5 Crystal Compensation With Periodic Synchronization

The periodic scheduling floods described in the previous section are initiated at a fixed interval. This, combined with the time-of-arrival estimates provided by the DW1000 (or similar UWB transceiver) provide an accurate estimation of the crystal offset between the master node and all other nodes participating in the Glossy flood.

This offset allows for the compensation of the crystal frequency used on each node. One method to perform this crystal offset compensation is through the fine-tuning of the

crystal's loaded capacitance. The DW1000 provides a method to fine-tune this capacitance at increments of 0.25 pF. This, combined with a careful selection of external loading capacitance depending on the crystal accuracy specification, can synchronization accuracy of 0.5 ppm between all nodes participating in the time synchronization and scheduling operations.

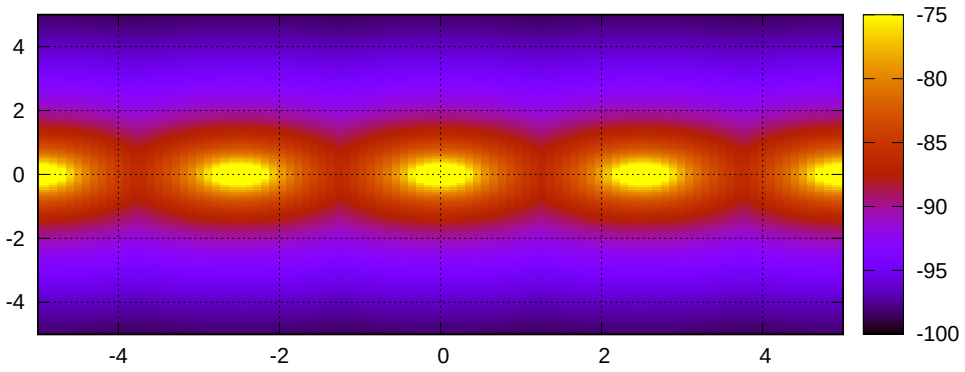

(a) Non-Simultaneous Transmissions

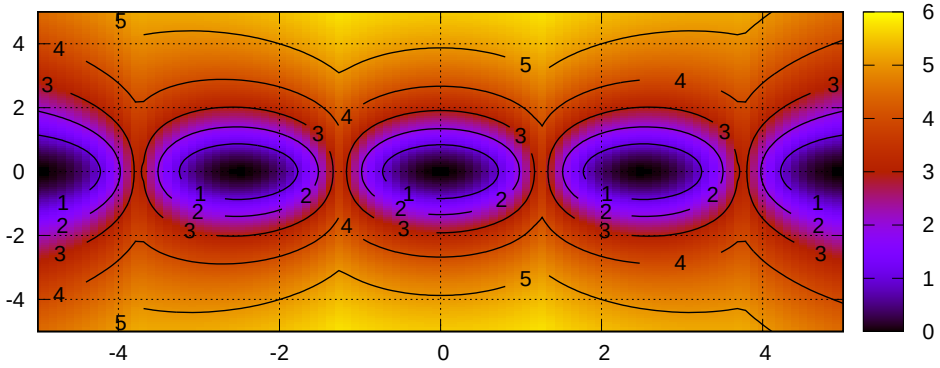

(b) Simultaneous Transmissions (dB increase in RX signal power)

Figure 6.7: **Combined Signal Power vs. Receiver Position for UWB Floods.** Compares the interference potential of five non-simultaneous transmitters (no flood) vs. multiple simultaneous transmitters (flooding). Only portions of the environment which are far away from the anchors experience any significant increase in signal power (>3 dB), where UWB signal power is low to begin with. A careful evaluation of the effects of UWB flooding on narrowband radios should be performed, considering the alternate of prolonged, per-tag sequential floods.

6.6 Interference Potential With UWB Floods

UWB floods increase the channel power over that attained with only one node transmitting at a time. This means that there is greater potential for unintended disruption to other radios occupying the same bands. In this section, we will detail the total impact on channel power and evaluate whether Glossy UWB floods will have any significant impact on other technologies occupying the same space.

Figure 6.7 shows a mapping of the maximum signal strength obtainable at a 3.1 GHz frequency across a floor plan 10x10 m^2 in size. Five nodes are placed along the X axis. The first figure shows the signal strength attainable across the space without constructive interference). The second figure shows the difference in signal strength attainable with constructive interference (all nodes transmitting simultaneously). Minimal changes in signal strength are seen close to each node, where narrowband interference is most likely, while a significant increase in receive power is observed further away from the array. Future work should detail placement requirements and perform a more thorough evaluation of the potential effects of UWB flooding on existing narrowband technologies.

6.7 Limits of Time Synchronization Accuracy

Time synchronization has been an invaluable byproduct of the GPS network for its use in providing approximately 100 nanoseconds in time synchronization accuracy to anywhere on Earth. This allows for a variety of applications ranging from the synchronization of national power transmission to the time stamping of business transactions. The techniques discussed in this section may be able to be leveraged to provide similar levels of accuracy in indoor environments, but a bounds on this accuracy will need to be explored in future work.

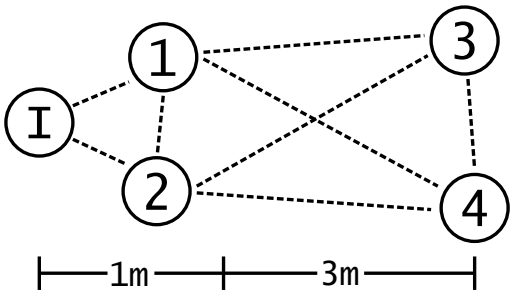

(a) Node Layout

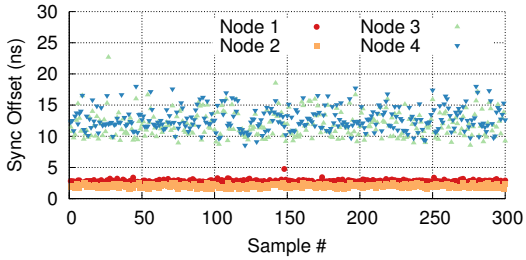

(b) Flood Timing Synchronization

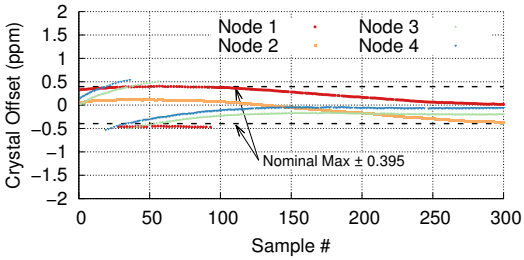

(c) Flooding-Based Crystal Synchronization

Figure 6.8: **Flooding Synchronization.** We lay out a small network of attenuated nodes with single-hop connectivity shown in (a). For UWB packets to constructively interfere, transmitter synchronization accuracy must exceed the position modulation timing granularity (64.10 ns for 6.8 Mbps transmission). Synchronization across the entire flooding network is affected by the accuracy of transmit and receive timestamps and the cumulative flood traversal time (i.e. the flooding protocol does not compensate for time-of-flight). In (b) we measure the offset of each node from the flood initiator. Average time synchronization is well within the required accuracy, yet those nodes situated after the first hop exhibit higher synchronization noise due to limitations in the transmit timing granularity of the DW1000 transceiver. In (c), we show how SurePoint adjusts the crystal frequency tuning to minimize cumulative timing error and enable reliable flooding over multiple hops. With a 23 ppm tuning range and 30 tuning steps, each node can (ideally) tune crystal frequency to a resolution of 0.79 ppm. SurePoint nodes compare the measured and expected time between successive floods to adjust their crystal frequency to that of the master node. Crystal non-linearities and manufacturing variability cause the offset to briefly exceed $\pm^{0.79}/_{2}$ and oscillate between steps when near the quantization boundary.

6.8 Is Global Time Sync Too Inclusive?

The methods for UWB floods and scheduling which have been discussed so far have assumed a global schedule for all tags participating in the localization network. In some cases where the network covers many distinct areas, a global time schedule may be too limiting. Instead, those nodes which are within range of disparate sets of anchors should be able to occupy the same time slot, as no potentially disruptive interference will result. Each disparate set could use a different preamble code in order to minimize any bleed-through of weak signals into the recorded impulse response of tags operating together in the same time slot.

6.9 Summary

In this chapter, we have identified many of the pain points which occur with the realization of localization systems deployed in indoor environments. These pain points include the deployment challenges of infrastructure, including time synchronization, the scheduling of transient tags, and the ability and reliability of supporting network-wide message capabilities.

Solutions to these challenges have been evaluated, including the introduction of the ultra-wideband flooding primitive which is able to achieve nanosecond-level time synchronization as well as improve the reliability of message transmission throughout the SurePoint network. Time synchronization has the potential to alleviate the need for hard-wired time synchronization usually used in time difference-of-arrival systems. Global communication throughout the network with wideband floods requires no additional network wiring for infrastructure nodes. Furthermore, these techniques enable the scheduling of tags which dynamically enter and leave the network, as would be expected in a person navigation environment.

The deployment of infrastructure is a major challenge for the retrofit of localization technologies in pre-existing buildings. If it is to be performed, it requires significant effort, and a desire is placed on the ease and cost of deployment, as well as the flexibility of the infrastructure being deployed. Throughout this chapter and dissertation, we have attempted to reduce the burden of infrastructure deployment.

CHAPTER 7

Conclusion

Historically, the key to adapting a new technology lies in identifying its killer app. In this dissertation, we explore the realization of indoor localization spanning the applications of personal navigation, robotic navigation, and asset tracking, all of which are motivating examples exposing the widely-ranging performance requirements encountered in typical localization system design. Due to the coverage and accuracy afforded by ultra-wideband RF transmissions, we argue that UWB technologies will allow for a generic and low-cost solution to this wide-ranging domain. A number of tag designs are proposed along with a generic infrastructure architecture designed to support the different sensing modalities required by each application.

We first detail the design of the provisioned SurePoint tag, a design which leverages commodity UWB transceivers and antenna diversity to provide position estimates with a focus on providing accurate and responsive location as a service. The SurePoint design is evaluated with respect to its ability to provide personal navigation in difficult indoor environments where high accuracy, low latency, and a user-friendly API are favored at the cost of tag power and higher tag cost.

Harmonium leverages antenna diversity to achieve similar accuracy as SurePoint but separates the UWB transmit and receive functions to alleviate power and implementation complexity requirements at the tag. Harmonium tags are transmit-only, shifting UWB receive complexity to the deployed infrastructure. The Harmonium receivers consist of a bandstitched narrowband frontend – a much simpler architecture than prior UWB receive architectures – further enhancing its integrability into current infrastructure designs.

Slocalization combines the Harmonium bandstitched receiver with a bandstitched transmitter to achieve low-noise UWB channel sounding. This enables the Slocalization tag, an architecture which achieves orders of magnitude improvement in tag power draw at the cost of slower tag update rate. The application of Slocalization to asset tracking is considered due to the long periods of stationarity characteristic of many asset tracking scenarios.

Finally, issues of infrastructure deployment such as global communication, time synchronization, and scheduling are explored, lowering the barrier to adoption for the techniques and technologies discussed in this dissertation. With a careful understanding of the various methods described for determining location, we believe that UWB techniques will allow for the quick adoption of indoor tracking applications which have seen little commercial utilization in the past.

7.1 Future Work

Indoor localization system design has and will continue to be an active area of research interest due to its wide-ranging application potential. The ideas and methods discussed in this dissertation have focused on methods for generating, manipulating, and measuring the ultra-wideband channel for the purposes of indoor localization. Nevertheless, there are a large number of directions that deserve future study in refining the size, weight, accuracy, power, and deployability of these systems. The following sections describe some of the immediate areas for improvement which have been identified in the systems described so far.

7.1.1 SurePoint Enhancements

The current SurePoint protocol assumes equal localization update requirements across all tags in the network. However, there may be times in which certain nodes may be under strict power requirements or may not be moving enough to warrant a quick update rate. Future implementations may be able to implement a more advanced scheduler by appending an update request frequency to their schedule request packet when requesting to join the network.

Motivated by this point, certain effort could be dedicated to improving the power consumption aspects of the ranging protocol's current implementation to fit a much larger number of application domains. First, the tags currently are used to participate (receive and flood) during the schedule and contention rounds. If the tag is under a strict power budget, the tag could skip flooding any messages during these times and only turn on to catch the periodic synchronization round messages to keep synchronized with the network.

Another avenue for improvement is applications in which the anchors span multiple distinct physical rooms or indoor spaces. In this case, a global schedule may not be optimal, as the broadcasts from a tag in a distant room will not likely collide. A means of dividing buildings into distinct RF zones with concurrent tag transmission could increase utilization.

Another interesting avenue to explore is feedback from the anchor network, which could detect when tags are moving through the environment and automatically hand it off between distinct RF zones.

7.1.2 Making Slocalization Less Slow

While Slocalization's current performance is acceptable for a large array of devices and applications, there are numerous improvements that could accelerate and improve localization performance. First, the RF frontends used exhibit an approximately 12 dB noise figure across the range of utilized frequencies. This yields the potential for improvement with the addition of a low-noise amplifier at each anchor receive antenna.

Currently, Slocalization uses omnidirectional antennas for maximum flexibility of anchor placement. In the WiTrack system, directional antennas are employed following the argument that the most likely deployment scenario is "in the walls". The same is likely true for Slocalization in many cases and the system could see similar gains. Further including system-level knowledge such as constraining tag location to "inside the room" could also improve localization performance and help eliminate outliers.

The instantaneous bandwidth is smaller than that attainable with the hardware utilized, as the gigabit Ethernet communication used by the USRP N210 bottlenecks throughput. Larger instantaneous bandwidth could be attainable by leveraging averaging on the FPGA fabric, lowering the necessary Ethernet bandwidth and therefore increasing the sweep rate and attainable update rate given the same specifications.

7.1.3 Slocalizing Smart Dust

Recently, there has been a growing interest in and several demonstrations of viable *millimeter*-scale systems [54]. Fundamentally, a Slocalization tag requires very little: a stable clock source, a shift register, and a variable impedance antenna element. Leveraging recent advances in near threshold circuit and oscillator designs, these components could be realized with an energy budget on the order of nanoamps [12]. As nodes shrink, however, their physical antennas necessarily shrink as well, significantly reducing antenna gain. Electrically small UWB antennas are still an active area of research, but the smallest antennas yielding high efficiency (near 0 dBi) are around 1 cm in diameter [97]. A recent effort to optimize antennas for mm-scale nodes has shown that narrowband mm-scale antennas realize gains of around -15 dBi within the Slocalization frequency range [11]. Assuming a similar correlation to achievable UWB antenna gain along with the doubling in path loss affected by the backscatter link, the the Slocalization system would be required to realize

another 30 dB of gain. This 30 dB of additional gain makes the integration times required for the current system intractable, but higher instantaneous bandwidth (up to $49\times = 17$ dB) and lower noise figure (12 dB) would be able to almost completely make up for these shortcomings.

BIBLIOGRAPHY

[1] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller. 3D tracking via body radio reflections. In *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation*, NSDI'14, pages 317–329. USENIX Association, 2014.

[2] F. Adib, H. Mao, Z. Kabelac, D. Katabi, and R. C. Miller. Smart homes that monitor breathing and heart rate. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 837–846, New York, NY, USA, 2015. ACM.

[3] R. Azim, M. T. Islam, and N. Misran. Compact tapered-shape slot antenna for UWB applications. *Antennas and Wireless Propagation Letters, IEEE*, 10:1190–1193, 2011.

[4] Bitcraze. The Crazyflie Nano Quadcopter.

[5] R. Blazquez, F. S. Lee, D. D. Wentzloff, P. P. Newaskar, J. D. Powell, and A. P. Chandrakasan. Digital architecture for an ultra-wideband radio receiver. In *Vehicular Technology Conference*, volume 2, pages 1303–1307, 2003.

[6] P. Bolliger. Redpin – adaptive, zero-configuration indoor localization through user collaboration. In *Proceedings of the First ACM International Workshop on Mobile Entity Localization and Tracking in GPS-less Environments*, MELT'08, pages 55–60. ACM, 2008.

[7] M. S. Brandstein, J. E. Adcock, and H. F. Silverman. A closed-form location estimator for use with room environment microphone arrays. *Speech and Audio Processing, IEEE Transactions on*, 5(1):45–50, 1997.

[8] R. Bucher and D. Misra. A synthesizable VHDL model of the exact solution for 3D hyperbolic positioning system. *VLSI Design*, 15, 2002.

[9] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.

[10] Y. Chan and K. Ho. A simple and efficient estimator for hyperbolic location. *Signal Processing, IEEE Transactions on*, 42(8):1905–1915, 1994.

[11] Y. Chen, N. Chiotellis, L. X. Chuo, C. Pfeiffer, Y. Shi, R. G. Dreslinski, A. Grbic, T. Mudge, D. D. Wentzloff, D. Blaauw, and H. S. Kim. Energy-autonomous wireless communication for millimeter-scale internet-of-things sensor nodes. *IEEE Journal on Selected Areas in Communications*, 34(12):3962–3977, Dec 2016.

[12] M. Choi, T. Jang, S. Bang, Y. Shi, D. Blaauw, and D. Sylvester. A 110 nw resistive frequency locked on-chip oscillator with 34.3 ppm/°C temperature stability for system-on-chip designs. *IEEE Journal of Solid-State Circuits*, 51(9):2106–2118, Sept 2016.

[13] G. Conte, M. De Marchi, A. A. Nacci, V. Rana, and D. Sciuto. BlueSentinel: A first approach using iBeacon for an energy efficient occupancy detection system. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, BuildSys'14, pages 11–19, 2014.

[14] D. Dardari, R. D'Errico, C. Roblin, A. Sibille, and M. Z. Win. Ultrawide bandwidth rfid: The next generation? *Proceedings of the IEEE*, 98(9):1570–1582, Sept 2010.

[15] G. Deak, K. Curran, and J. Condell. A survey of active and passive indoor localisation systems. *Computer Communications*, 35(16):1939–1954, 2012.

[16] S. DeBruin, B. Campbell, and P. Dutta. Monjolo: An energy-harvesting energy meter architecture. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, SenSys '13, pages 18:1–18:14, New York, NY, USA, 2013. ACM.

[17] DecaWave. Antenna delay calibration of DW1000-based products and systems.

[18] DecaWave. ScenSor DW1000. http://www.decawave.com/.

[19] R. D'Errico, M. Bottazzi, F. Natali, E. Savioli, S. Bartoletti, A. Conti, D. Dardari, N. Decarli, F. Guidi, F. Dehmas, et al. An uwb-uhf semi-passive rfid system for localization and tracking applications. In *RFID-Technologies and Applications (RFID-TA), 2012 IEEE International Conference on*, pages 18–23. IEEE, 2012.

[20] E. Dhalgren and H. Mahmood. Evaluation of indoor positioning based on Bluetooth Smart technology. Master's thesis, Chalmers, Jun 2014.

[21] T. Domain. http://timedomain.com.

[22] P. Dutta, R. Musaloiu-E., I. Stoica, and A. Terzis. Wireless ack collisions not considered harmful. In *Proceedings of the Seventh Workshop on Hot Topics in Networks*, HotNets-VII, Oct 2008.

[23] Ettus Research LLC. USRP1 Bus Series, Ettus Research LLC. https://www.ettus.com/product/category/USRP-Bus-Series.

[24] C. Falsi, D. Dardari, L. Mucchi, and M. Z. Win. Time of arrival estimation for uwb localizers in realistic environments. *EURASIP Journal on Advances in Signal Processing*, 2006(1):032082, 2006.

[25] Federal Communications Commission. First report and order 02-48. Technical report, Federal Communications Commission, February 2002.

[26] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele. Low-power wireless bus. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, pages 1–14. ACM, 2012.

[27] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient network flooding and time synchronization with glossy. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 73–84. IEEE, 2011.

[28] M. Gerding, T. Musch, and B. Schiek. Generation of short electrical pulses based on bipolar transistors. *Advances in Radio Science*, 2(1):7–12, 2005.

[29] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: A look at positioning aspects for future sensor networks. *IEEE Signal Processing Magazine*, 22(4):70–84, July 2005.

[30] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: A look at positioning aspects for future sensor networks. *IEEE Signal Processing Magazine*, 22(4):70–84, July 2005.

[31] F. Gustafsson and F. Gunnarsson. Positioning using time-difference of arrival measurements. In *Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03). 2003 IEEE International Conference on*, volume 6, pages VI–553–6 vol.6, April 2003.

[32] J. Han and C. Nguyen. A new ultra-wideband, ultra-short monocycle pulse generator with reduced ringing. *Microwave and Wireless Components Letters, IEEE*, 12(6):206–208, June 2002.

[33] S. Hantscher, A. Reisenzahn, H. Kainmüller, and C. G. Diskus. Hardware concepts for the sequential sampling of repetitive pulse radar echoes in cost-efficient ultra-wideband transceivers. *Microwave and Optical Technology Letters*, 52(3):585–591, 2010.

[34] V. Heiries, K. Belmkaddem, F. Dehmas, B. Denis, L. Ouvry, and R. D'Errico. Uwb backscattering system for passive rfid tag ranging and tracking. In *Ultra-Wideband (ICUWB), 2011 IEEE International Conference on*, pages 489–493. IEEE, 2011.

[35] W. Hirt. The european uwb radio regulatory and standards framework: Overview and implications. In *2007 IEEE International Conference on Ultra-Wideband*, pages 733–738, Singapore, Sept 2007. IEEE.

[36] S. Hoyos and B. M. Sadler. Ultra-wideband analog-to-digital conversion via signal expansion. *Vehicular Technology, IEEE Transactions on*, 54:1609–1622, 2005.

[37] P. Hu, P. Zhang, and D. Ganesan. Laissez-faire: Fully asymmetric backscatter communication. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, SIGCOMM '15, pages 255–267, New York, NY, USA, 2015. ACM.

[38] D. Humphrey and M. Hedley. Super-resolution time of arrival for indoor localization. In *IEEE International Conference on Communications*, 2008.

[39] IEEE Computer Society. Part 15.4: Low-rate wireless personal area networks (LR-WPANs). https://standards.ieee.org/getieee802/download/802.15.4-2011.pdf.

[40] E. Inc. Ekahau rtls. http://www.ekahau.com/real-time-location-system, June 2016.

[41] ISA. Special section: Wireless, tracking wireless. https://www.isa.org/standards-and-publications/isa-publications/intech-magazine/2008/april/special-section-wireless-tracking-wireless/, Apr. 2008.

[42] B. Kempke, P. Pannuto, B. Campbell, and P. Dutta. SurePoint: Exploiting ultra wideband flooding and diversity to provide robust, scalable, high-fidelity indoor localization. In *Proceedings of the 14th ACM Conference on Embedded Networked Sensor Systems*, SenSys'16, pages 137–149, New York, NY, USA, November 2016. ACM.

[43] B. Kempke, P. Pannuto, and P. Dutta. Harmonia: Wideband spreading for accurate indoor RF localization. In *2014 ACM Workshop on Hot Topics in Wireless*, HotWireless '14, September 2014.

[44] B. Kempke, P. Pannuto, and P. Dutta. PolyPoint: Guiding indoor quadrotors with ultra-wideband localization. In *2015 ACM Workshop on Hot Topics in Wireless*, HotWireless '15, September 2015.

[45] B. Kempke, P. Pannuto, and P. Dutta. Harmonium: Asymmetric, bandstitched UWB for fast, accurate, and robust indoor localization. In *Proceedings of the 15th International Conference on Information Processing in Sensor Networks*, IPSN'16, pages 15:1–15:12, Piscataway, NJ, USA, April 2016. IEEE Press.

[46] K. Klipp, J. Willaredt, H. Rosé, and I. Radusch. Low cost high precision indoor localization system fusing intertial and magnetic field sensor data with radio beacons. *Second Annual Microsoft Indoor Localization Competition*, 2015.

[47] S. Kumar, S. Gil, D. Katabi, and D. Rus. Accurate indoor localization with zero start-up cost. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom'14, pages 483–494. ACM, 2014.

[48] Y.-S. Kuo, P. Pannuto, K.-J. Hsiao, and P. Dutta. Luxapose: Indoor positioning with mobile phones and visible light. In *The 20th Annual International Conference on Mobile Computing and Networking*, MobiCom'14, September 2014.

[49] D. Lachartre, B. Denis, D. Morche, L. Ouvry, M. Pezzin, B. Piaget, J. Prouvée, and P. Vincent. A 1.1 nJ/b 802.15.4a-compliant fully integrated UWB transceiver in 0.13 µm CMOS. In *Solid-State Circuits Conference*, pages 312–313, 2009.

[50] S. Lanzisera, D. Zats, and K. S. Pister. Radio frequency time-of-flight distance measurement for low-cost wireless sensor localization. *IEEE Sensors Journal*, 11(3):837–845, 2011.

[51] P. Lazik and A. Rowe. Indoor pseudo-ranging of mobile devices using ultrasonic chirps. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, SenSys'12, pages 99–112. ACM, 2012.

[52] J.-S. Lee and C. Nguyen. Novel low-cost ultra-wideband, ultra-short-pulse transmitter with MESFET impulse-shaping circuitry for reduced distortion and improved pulse repetition rate. *Microwave and Wireless Components Letters, IEEE*, 11(5):208–210, May 2001.

[53] J.-Y. Lee and R. A. Scholtz. Ranging in a dense multipath environment using an UWB radio link. *IEEE Journal on Selected Areas in Communications*, 20(9):1677–1683, Dec 2002.

[54] Y. Lee, S. Bang, I. Lee, Y. Kim, G. Kim, M. H. Ghaed, P. Pannuto, P. Dutta, D. Sylvester, and D. Blaauw. A modular 1 mm^3 die-stacked sensing platform with low power I^2C inter-die communication and multi-modal energy harvesting. *IEEE Journal of Solid-State Circuits*, 48(1):229–243, 2013.

[55] E. Leitinger, P. Meissner, C. Rüdisser, G. Dumphart, and K. Witrisal. Evaluation of position-related information in multipath components for indoor positioning. *IEEE Journal on Selected Areas in communications*, 33(11):2313–2328, 2015.

[56] H. Liu, H. Darabi, P. Banerjee, and J. Liu. Survey of wireless indoor positioning techniques and systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37(6):1067–1080, 2007.

[57] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith. Ambient backscatter: Wireless communication out of thin air. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, pages 39–50, New York, NY, USA, 2013. ACM.

[58] I. Lu, N. Weste, and S. Parameswaran. ADC precision requirement for digital ultra-wideband receivers with sublinear front-ends: A power and performance perspective. In *VLSI Design*, 2006.

[59] D. Lymberopoulos, J. Liu, X. Yang, A. Naguib, A. Rowe, N. Trigoni, and N. Moayeri. Microsoft indoor localization competition – IPSN 2015.

[60] D. Lymberopoulos, J. Liu, Y. Zhang, P. Dutta, Y. Xue, and A. Rowe. Microsoft indoor localization competition – IPSN 2016.

[61] R. Mautz. Indoor positioning technologies, 2012.

[62] A. F. Molisch, K. Balakrishnan, D. Cassioli, C.-C. Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak. IEEE 802.15.4a channel model - final report, 2004.

[63] A. F. Molisch, D. Cassioli, C. C. Chong, S. Emami, A. Fort, B. Kannan, J. Karedal, J. Kunisch, H. G. Schantz, K. Siwiak, and M. Z. Win. A comprehensive standardized model for ultrawideband propagation channels. *IEEE Transactions on Antennas and Propagation*, 54(11):3151–3166, Nov 2006.

[64] Y. Mulgaonkar, M. Whitzer, B. Morgan, C. M. Kroninger, A. M. Harrington, and V. Kumar. Power and weight considerations in small, agile quadrotors. In *SPIE Defense+Security*, pages 90831Q–90831Q, 2014.

[65] NaturalPoint. OptiTrack.

[66] 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*, MobiCom'00, pages 32–43. ACM, 2000.

[67] A. Rabbachin, T. Q. Quek, P. C. Pinto, I. Oppermann, and M. Z. Win. Uwb energy detection in the presence of multiple narrowband interferers. In *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, pages 857–862. IEEE, 2007.

[68] N. Rajagopal, P. Lazik, and A. Rowe. Visual light landmarks for mobile devices. In *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*, IPSN'14, pages 249–260. IEEE Press, 2014.

[69] R. M. Richardson. Remotely actuated radio frequency powered devices, July 1963. US Patent 3,098,971.

[70] Z. Sahinoglu, S. Gezici, and I. Guvenc. Ultra-wideband positioning systems. *Cambridge, New York*, 2008.

[71] T. Sathyan, D. Humphrey, and M. Hedley. WASP: A system and algorithms for accurate radio localization using low-cost hardware. *IEEE Transactions on Systems, Man, and Cybernetics – Part C*, 41(2):211–222, Mar. 2011.

[72] T. Sathyan, D. Humphrey, and M. Hedley. WASP: A system and algorithms for accurate radio localization using low-cost hardware. *IEEE Transactions on Systems, Man, and Cybernetics – Part C*, 41(2), Mar. 2011.

[73] M. Segura, H. Hashemi, C. Sisterna, and V. Mut. Experimental demonstration of self-localized ultra wideband indoor mobile robot navigation system. In *Indoor Positioning and Indoor Navigation (IPIN), 2010 International Conference on*, pages 1–9. IEEE, 2010.

[74] STMicroelectronics. STM32F031G6U6. http://www.st.com/en/microcontrollers/stm32f031g6.html.

[75] M. Sugano, T. Kawazoe, Y. Ohta, and M. Murata. Indoor localization system using rssi measurement of wireless sensor network based on zigbee standard. *Target*, 538:050, 2006.

[76] L. Sun, S. Sen, D. Koutsonikolas, and K.-H. Kim. Widraw: Enabling hands-free drawing in the air on commodity wifi devices. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 77–89. ACM, 2015.

[77] V. Talla and J. R. Smith. Hybrid analog-digital backscatter: A new approach for battery-free sensing. In *2013 IEEE International Conference on RFID (RFID)*, pages 74–81, Orlando, FL, USA, April 2013. IEEE.

[78] S. P. Tarzia, P. A. Dinda, R. P. Dick, and G. Memik. Indoor localization without infrastructure using the acoustic background spectrum. In *Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services*, MobiSys'11, pages 155–168. ACM, 2011.

[79] D. Z. Thai, M. Trinkle, A. Hashemi-Sakhtsari, and T. Pattison. Speaker localisation using time difference of arrival. Technical report, DTIC Document, 2008.

[80] K. M. S. Thotahewa, J.-M. Redouté, and M. R. Yuce. Hardware architectures for IR-UWB-based transceivers. In *Ultra Wideband Wireless Body Area Networks*, pages 67–81. Springer, 2014.

[81] Ubisense. Series 7000 Compact Tag. http://www.ubisense.net/en/media/pdfs/products_pdf/uk/80553_series_7000_compact_tag.pdf.

[82] D. Vasisht, S. Kumar, and D. Katabi. Decimeter-level localization with a single WiFi access point. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 165–178, Santa Clara, CA, Mar. 2016. USENIX Association.

[83] M. Verhelst, N. Van Helleputte, G. Gielen, and W. Dehaene. A reconfigurable, 0.13 µm CMOS 110 pJ/pulse, fully integrated IR-UWB receiver for communication and sub-cm ranging. In *Solid-State Circuits Conference*, pages 250–251, 2009.

[84] Vicon. T-Series: Precise Motion Capture in Any Volume Anywhere.

[85] C. Walton. Portable radio frequency emitting identifier, May 1983. US Patent 4,384,288.

[86] E. A. Wan and A. S. Paul. A tag-free solution to unobtrusive indoor tracking using wall-mounted ultrasonic transducers. In *Indoor Positioning and Indoor Navigation (IPIN), 2010 International Conference on*, pages 1–10. IEEE, 2010.

[87] J. Wang, D. Vasisht, and D. Katabi. RF-IDraw: Virtual touch screen in the air using RF signals. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM'14, pages 235–246. ACM, 2014.

[88] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, 1997.

[89] D. Wentzloff and A. Chandrakasan. A 3.1-10.6 GHz ultra-wideband pulse-shaping mixer. In *Radio Freq. Integrated Circuits (RFIC) Symposium*, pages 83–86, June 2005.

[90] K. Witrisal, E. Leitinger, S. Hinteregger, and P. Meissner. Bandwidth scaling and diversity gain for ranging and positioning in dense multipath channels. *IEEE Wireless Communications Letters*, 5(4):396–399, 2016.

[91] K. Witrisal, P. Meissner, E. Leitinger, Y. Shen, C. Gustafson, F. Tufvesson, K. Haneda, D. Dardari, A. F. Molisch, A. Conti, et al. High-accuracy localization for assisted living: 5g systems will turn multipath channels from foe to friend. *IEEE Signal Processing Magazine*, 33(2):59–70, 2016.

[92] J. Wright and S. Nocedal. *Numerical Optimization*. SciPy, 1999.

[93] K. Wu, J. Xiao, Y. Yi, M. Gao, and L. M. Ni. FILA: Fine-grained indoor localization. In *INFOCOM, 2012 Proceedings IEEE*, pages 2210–2218. IEEE, 2012.

[94] J. Xiong and K. Jamieson. ArrayTrack: A fine-grained indoor location system. In *Usenix NSDI*, 2013.

[95] D. Yang, A. E. Fathy, H. Li, M. Mahfouz, and G. D. Peterson. Millimeter accuracy UWB positioning system using sequential sub-sampler and time difference estimation algorithm. In *Radio and Wireless Symposium (RWS), 2010 IEEE*, pages 539–542, 2010.

[96] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu. Tagoram: Real-time tracking of mobile RFID tags to high precision using COTS devices. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom'14, pages 237–248. ACM, 2014.

[97] T. Yang, W. A. Davis, and W. L. Stutzman. The design of ultra-wideband antennas with performance close to the fundamental limit. In *Proc. URSI General Assembly*, page 234, 2008.

[98] P. Zhang, J. Gummeson, and D. Ganesan. Blink: A high throughput link layer for backscatter communication. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*, MobiSys '12, pages 99–112, New York, NY, USA, 2012. ACM.

[99] Y. Zheng, M. A. Arasu, K.-W. Wong, Y. J. The, A. P. H. Suan, D. D. Tran, W. G. Yeoh, and D.-L. Kwong. A 0.18 µm CMOS 802.15.4a UWB transceiver for communication and localization. In *Solid-State Circuits Conference*, pages 118–600, 2008.

APPENDIX A

Design Resources

Design resources for the projects covered in this dissertation are all freely available online from the Lab11 website. Specific code repositories for these projects are listed below.

SurePoint Personal Localization System

`http://github.com/lab11/polypoint`

Harmonium Robotic Navigation System

`http://github.com/lab11/fast-square`

Slocalization Asset Tracking System

`http://github.com/lab11/slocalization`