SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization

Benjamin Kempke, Pat Pannuto, Bradford Campbell, and Prabal Dutta
Electrical Engineering and Computer Science Department
University of Michigan, Ann Arbor, MI 48109
{bpkempke,ppannuto,bradjc,prabal}@umich.edu

ABSTRACT

We present SurePoint, a system for drop-in, high-fidelity indoor localization. SurePoint builds on recently available commercial ultra-wideband radio hardware. While ultra-wideband radio hardware can provide the timing primitives necessary for a simple adaptation of two-way ranging, we show that with the addition of frequency and spatial diversity, we can achieve a 53% decrease in median ranging error. Because this extra diversity requires many additional packets for each range estimate, we next develop an efficient broadcast ranging protocol for localization that ameliorates this overhead. We evaluate the performance of this ranging protocol in stationary and fast-moving environments and find that it achieves up to 0.08 m median error and 0.53 m 99th percentile error. As ranging requires the tag to have exclusive access to the channel, we next develop a protocol to coordinate the localization of multiple tags in space. This protocol builds on recent work exploiting the constructive interference phenomenon. The ultra-wideband PHY uses a different modulation scheme compared to the narrowband PHY used by previous work, thus we first explore the viability and performance of constructive interference with ultra-wideband radios. Finally, as the ranging protocol requires careful management of the ultra-wideband radio and tight timing, we develop TriPoint, a dedicated “drop-in” ranging module that provides a simple I²C interface. We show that this additional microcontroller demands only marginal energy overhead while facilitating interoperability by freeing the primary microcontroller to handle other tasks.

CCS Concepts

•Computer systems organization → Embedded systems; •Networks → Location based services; Application layer protocols; •Hardware → Communication hardware, interfaces and storage;

Keywords

Indoor Localization, Ultra-Wideband Communications, Constructive Interference

This work is licensed under a Creative Commons Attribution International 4.0 License.

SenSys '16 November 14-16, 2016, Stanford, CA, USA

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4263-6/16/11.

DOI: <http://dx.doi.org/10.1145/2994551.2994570>

1. INTRODUCTION

Indoor localization is a well-studied and well-motivated problem. Over the last decade, researchers have leveraged GSM [5], WiFi [23, 39], Bluetooth [7, 10], ultra-wideband RF [2, 35, 38], acoustics [27, 37], magnetics [22], LIDAR [34], visible light communication [24, 33], power line communication [32], and other technologies to devise more accurate, faster, and more economical (in terms of energy and dollars) localization systems.

The research community has made great strides by designing indoor localization systems using multiple modalities that achieve decimeter-accurate localization quality (e.g. ALPS [26], Harmonium [21], Chronos [39], and Luxapose [24]). Given the growing abundance of systems that offer very high accuracy, perhaps continuing to evaluate systems on accuracy alone is no longer useful, and indeed decimeter level error is good enough. As localization technology begins to mature, we argue that it is time to focus on the other aspects necessary for a viable and reliable localization system.

To that end, we present SurePoint, a decimeter-accurate, robust, reliable, scalable, and easily integrable localization system. SurePoint is able to achieve a 29 cm median position error and 99 percent of ranging error is *within 53 cm*. SurePoint builds on our previous system, PolyPoint [20], but adds robustness, higher accuracy, an improved location solver, and support for localizing multiple tags. This entire system is then packaged into a drop-in localization module that can be added to other hardware designs.

SurePoint achieves its high accuracy by leveraging the high-fidelity time-of-arrival primitive commonly provided by ultra-wideband radios. SurePoint further improves on this raw primitive with the addition of frequency and polarization diversity by exploiting multiple communication channels and multiple antennas.

Supporting multiple tags in RF-based systems is often overlooked. While unidirectional broadcast systems, such as ALPS [26] or Luxapose [24] naturally support an unbounded number of devices, recent RF-based systems such as WiTrack [3] or Harmonium [21] do not address the challenge of supporting multiple tags or require burdensome hardware to coordinate tightly-synchronized transmissions. SurePoint takes inspiration from the recent Low Power Wireless Bus [14] to build a ranging protocol that dynamically adapts to the number of tags currently in the environment.

Using ultra-wideband, while advantageous for ranging, presents an interesting challenge when coupled with a design similar to the Low Power Wireless Bus. That bus protocol relies on Glossy floods [15] to synchronize nodes, and Glossy floods rely on the constructive interference phenomenon. Loosely, constructive interference occurs when two radios send the exact same packet at sufficiently the exact same time such that the collision of the two packets is additive, resulting in a more robust transmission. Previous work in constructive interference has focused on narrowband 802.15.4,

which uses a different modulation scheme than the UWB 802.15.4 PHY. We explore the constructive interference phenomenon in the UWB channel, demonstrating that it continues to improve transmission performance with UWB as well.

Finally, we observe that implementing a reliable and accurate localization system is challenging. Obtaining consistent location updates often requires meeting tight timing constraints, performing expensive processing, adapting to local conditions by selecting appropriate algorithmic parameters, and integrating this complexity with the system or device to be localized. For these reasons, we introduce TriPoint, a drop-in hardware module that provides localization as a simple I²C peripheral. We show how prefabricated modules help to address the calibration problem, ease system integration, and ultimately make locating any Bluetooth-enabled device, such as a smartphone or tablet, as simple as sticking a tag to it.

Contributions and Roadmap. This paper makes the following contributions:

- We study the impact of frequency and polarization diversity on the UWB ranging primitive in [Section 2](#). We find that for a modest increase in complexity, we can improve the median accuracy of the raw UWB range primitive 2x from 0.17 m to 0.08 m and the 99th percentile 2.5x from 1.30 m to 0.53 m.
- While diversity improves the raw performance of UWB ranging, it demands many additional packets for each range measurement. As indoor localization schemes require many ranges to many anchors in the environment, [Section 3](#) next explores how a single tag can efficiently and robustly collect ranges from all nearby anchors using a broadcast ranging protocol.
- We next aim to support the localization of multiple tags by leveraging recent work in constructive interference-based protocols [14]. Previous work, however, used narrowband 802.15.4 radios, which use O-QPSK modulation. In 802.15.4a, the UWB PHY uses BPM-BPSK. [Section 4](#) explores the viability and efficacy of constructive interference with the UWB 802.15.4 PHY, enabling its use for our protocol.
- Leveraging constructive interference and drawing inspiration from the Low Power Wireless Bus [14], we design a protocol for scaling localization support to multiple tags in the same environment. [Section 5](#) introduces the new protocol and demonstrates how SurePoint scales to handle multiple tags in the same environment.
- Finally in [Section 6](#) we introduce the TriPoint module, a drop-in hardware component to easily add SurePoint to any system.

We then evaluate each of these contributions in turn and close with a discussion of some of the limitations of our current implementation, some of the open avenues for further research, and comparison with some of the other state-of-the-art localization systems.

2. RANGING AND DIVERSITY

The SurePoint system leverages ultra-wideband (UWB) radios to capture high-fidelity range estimates between nodes. In this section, we provide a brief background of UWB and then explore methods to improve on the raw UWB ranging primitive, specifically frequency and polarization diversity.

2.1 Background

A wide breadth of RF localization research has shown that high fidelity characterizations of the indoor RF channel (and subsequently high localization accuracy) requires a large amount of bandwidth. For this reason, ultra-wideband techniques have shown superior performance in recent head-to-head comparisons of localization technologies [29, 30].

UWB radios spread RF power across a much wider bandwidth than traditional narrowband radios such as WiFi, 802.15.4, or RFID. For this reason, the amount of power that can be transmitted must be kept very low to avoid any destructive effects to narrowband radios occupying the same RF bands. With these limitations in mind, many international regulatory agencies [13] have recently enacted the unrestricted use of large portions of RF spectrum for UWB radios which abide by pre-defined power spectral density requirements.

A few commercial entities (Time Domain [2], UbiSense [38], Nanotron [1]) have capitalized on this opportunity through the development of separate, custom UWB RF localization systems which primarily target industrial, high price-point applications. Broader, consumer-grade localization adoption has found little traction due to the high per-tag and infrastructure costs of these systems, along with no concentrated efforts in protocol standardization. This barrier to adoption was upended through the addition of a UWB PHY to the 802.15.4a standard [18]. Subsequently, DecaWave released an 802.15.4a-compliant UWB radio, the DW1000, targeting a broader adoption of UWB radios for use in high-fidelity indoor localization.

2.2 Ranges in a UWB Channel

Most current localization systems use either sound, light, or radio wave propagation to determine an object’s physical location. Radio wave propagation – unlike both sound and light – has the ability to pass through many objects which may be opaque to light and/or sound. This provides two major benefits for localization systems which utilize RF. First, RF localization systems can provide better coverage in crowded settings. Secondly, any infrastructure can be deployed in a visually unobtrusive way, such as behind architectural elements or building elements such as drywall.

RF signals travel at high speed (3.0×10^8 m/sec), which makes their use in time-based systems difficult. Sampling requirements for meter-level resolution are extremely high ($\frac{3.0 \times 10^8 \text{ m/sec}}{1 \text{ m}} = 300 \text{ MHz}$), which makes receivers using off-the-shelf components costly and complex. Even 500 MHz of bandwidth—the typical UWB channel bandwidth—still has frequent overlap of closely-spaced multipath.

Time-based localization systems traditionally attempt to measure the time-of-arrival of the line-of-sight path. If measured accurately, this can be used to determine range or pseudo-range between two radios. Whenever a second path of propagation (i.e. a reflection) arrives within $\frac{c}{BW}$ seconds of the first path, distortion of the signal’s time-of-arrival determination can occur.

[Figure 1a](#) considers 1000 independent observations in an 802.15.4a channel model [31] to estimate the distribution of ranges seen. Range estimates are derived using the standard 20% height of the first peak in the channel impulse response. Notice that while most range errors are positive, occasionally, multipath fading will reduce the amplitude of the initial peak of the channel impulse response, causing the 20% height to return a premature estimate of the arrival time.

2.3 Leveraging RF Diversity

Given the unpredictable nature of the UWB channel, SurePoint recognizes that capturing a single channel estimate has the potential to introduce large amounts of variance in the range estimate. To combat this, SurePoint adds additional sources of diversity—multiple estimates of the UWB channel on a single node.

Figure 1: 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. [31], in (a) we take 100k samples of an active multipath environment and find that the estimated range between two nodes can vary by over a meter. Next, in (b) we collect 1.9 million samples while varying position, polarity, and channel in a real-world environment. From these analyses, assuming an unlimited number of independent observations, the 12th percentile will yield zero range error. The final plot, (c), gives the 12th percentile estimates from each position in the real-world samples, reducing the median error from 0.17 m to 0.08 m and 95th percentile errors from 0.59 m to 0.31 m.

Each SurePoint node features three physically separated antennas providing modest spatial diversity placed at 120° offsets to help combat nulls from cross-polarization of antennas. Furthermore, 802.15.4a has three different largely non-overlapping channels that enable SurePoint to exploit frequency diversity. All of the permutations of three tag antennas, three anchors antennas, and three channels results in 27 independent measurements of the UWB channel. In realistic environments, some of these observations will be dropped due to antenna nulls or destructive multipath interference. A protocol which is able to collect these many observations quickly and in a reliable manner is therefore required.

In Figure 1b we move a tag around an active space with numerous anchors and record all of the raw range estimates and then plot the error to each tag. From both the channel model and empirical analysis we derive that the 12th percentile estimate of independent samples minimizes the range error. The refinement of 12th percentile over the 10th percentile used in PolyPoint is informed by the new channel model study and a greater corpus of empirical range measurements.

Figure 1c takes the measured ranging data and computes the 12th percentile of each observation. The resulting distribution reduces the median error from 0.17 m to 0.08 m and the 95th percentile error from 0.59 m to 0.31 m.

Figure 2: Ranging Protocol. The SurePoint ranging protocol exploits frequency and spatial diversity with minimal impact on update rate. The protocol begins with 27 broadcast transmissions from the tag for each combination of tag antenna, anchor antenna, and RF channel. Next, the tag repeats the first packet on each channel (for robustness) so that the tag and anchor can compute their relative crystal offset. Finally, each anchor chooses a random time in the reply timeslot to respond to the tag with all of its observed measurements. The tag acknowledges the anchor reply if it reaches the tag. If an anchor does not receive an acknowledgment (e.g. packet dropped due to channel conditions or a collision), it retries on a new channel for up to three total tries.

3. FROM RANGING TO LOCALIZATION

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.

3.1 SurePoint Ranging Protocol (Figure 2)

Per the discussion in Section 2.3, SurePoint 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, PolyPoint developed the original broadcast ranging protocol which we describe here [20]. We improve upon this protocol by adding much-needed robustness. In our original PolyPoint 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.

SurePoint applies two optimizations to greatly reduce the number of transmissions required to integrate antenna and frequency diversity between SurePoint 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 PolyPoint 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 5](#).

3.2 Position Solver

In order to determine the tag's location, SurePoint 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. [16] and adds iterative point removal to a non-linear least squares solver [41] 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 [41]. Guided by Gezici's observations [16], we enhance the original PolyPoint solver by adding an iterative step that considers removing an anchor whenever the error fails to converge to better than 10 cm.

4. CONSTRUCTIVE INTERFERENCE

To efficiently localize multiple tags in the same space, we develop a time segmentation approach for SurePoint inspired by the Low-Power Wireless Bus scheduling scheme [14]. The prior work on low power wireless synchronization has shown that flooding-based protocols can achieve high synchronization accuracy with a minimal number of RF transmissions. These flooding-based time synchronization schemes (e.g. Glossy [15]) also benefit from high reliability due to the constructive interference effects of the multiple transmitters operating simultaneously during the periodic synchronization floods. SurePoint achieves a global time-slotted tag schedule by extending the flooding-based Low-Power Wireless Bus (LWB) architecture to use Glossy floods over 802.15.4a.

4.1 Constructive Interference in Theory

SurePoint uses the DecaWave DW1000 UWB transceiver [9]. The DW1000 provides the ability to measure packet reception time down to a precision of 15.65 *pico*seconds as well as the ability to schedule transmissions with a precision of 8 *nan*oseconds. This high precision is a primitive developed for achieving high RF localization accuracy, but has the additional benefit of enabling coordinated transmissions across nodes to an accuracy greater than the timing granularity of the UWB physical layer (64.10 ns during 6.8 Mbps data transmission). This is the first necessary condition to support Glossy floods.

The second necessary condition to support Glossy floods relies on the ability of the underlying channel to support constructive interference. Due to the black-box nature of the DW1000 receiver, we instead refer to the worst-case coherency condition to support constructive interference. This states that the two radio's local oscillators must not drift in phase by more than one half-cycle across an entire packet duration. In the case of our 264-bit 802.15.4a frames transmitted at 6.8 Mbps with a 64 μ s preamble, this equates to a coherency timing requirement of 103 μ s. With a 3.4944 GHz local oscillator, this equates to a necessary crystal accuracy of $0.5/3.4944 \times 10^9 / 103 \times 10^{-6} = 1.39$ ppm.

4.2 Constructive Interference in Practice

As our SurePoint implementation meets the stated requirements for constructive interference, we next perform a brief empirical analysis to verify that the constructive interference phenomenon operates as expected in the 802.15.4 UWB channel. In [Figure 4](#), we place three nodes in Faraday cages connected with variable attenuators. One node acts as an initiator and sends a packet which is echoed back precisely 10 ms later by one node, the other node,

Figure 3: 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.

Figure 4: Constructive Interference with UWB. To evaluate whether UWB exhibits constructive interference we place three nodes, an initiator and two echoers, into separate 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 802.15.4 UWB channel.

or both. In the case of both nodes replying, the packet reception rate by the initiator rises. This is strong evidence that constructive interference is occurring as we would expect the capture effect to achieve no better performance than a single tag on its own.

Constructive flooding protocols necessarily provide some degree of time synchronization across the network. Figure 3b explores the synchronization quality provided by SurePoint floods and finds that error is dominated by the time-of-flight distance from the flood initiator as the current SurePoint flooding protocol does not compensate for the distance between nodes.

Recall that our analysis of constructive interference for 802.15.4a UWB necessitated a crystal accuracy of at least 1.39 ppm. The DW1000 provides the ability to tune its crystal frequency via a 5-bit tuning register. Through the concurrent selection of a ± 10 ppm crystal along with appropriate loading capacitors to limit the tuning range, we can achieve sub-ppm crystal tuning resolution. The value of the desired crystal tuning is measured by determining the offset between the expected and actual time between floods. Figure 3c shows how SurePoint adapts the crystal offset over time to minimize synchronization error.

5. SUPPORTING MULTIPLE TAGS

The SurePoint system evaluated in this paper consists of a number of fixed-location nodes (anchors) along with one or more nodes to localize (tags). The anchors are assumed to be wall-powered and under no realistic power constraints, while the tags are assumed to be battery-powered, making their energy usage important.

For SurePoint, we are interested in the scheduling provided by the Low Power Wireless Bus (LWB), but not its data transmissions. We leverage the same scheduling and contention-based joining mechanisms as LWB, but then replace its data transmission with SurePoint ranging events. Because wall-powered nodes are unlikely to drop out of communication with the rest of the network, a number of simplifications can be made in SurePoint’s LWB-variant. One anchor node in the network is deemed to be the master and performs all tasks related to providing network time synchronization to all nodes as well as all scheduling for tags involved in the network.

5.1 Protocol Overview (Figure 5)

The master anchor’s purpose is two-fold in the SurePoint network. First, it starts a periodic synchronization flood every T_{round} (one second in our implementation) to keep a consistent timebase across all nodes in the network and kick off each localization round. Second, the master processes all incoming schedule requests from newly-seen tags and adds them to the periodic scheduling round.

The periodic synchronization floods aim to keep everyone involved in the SurePoint network synchronized to an accurate, global timebase. Those nodes who have missed one of the last two sync floods are assumed to be unsynchronized and are therefore disallowed from participating in any subsequent SurePoint operations until the time at which they have been deemed to be re-synchronized.

5.1.1 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

Figure 5: **Multi-Tag Protocol.** Every round begins with a scheduling flood initiated by the master anchor. This schedule includes the round duration which includes a contention interval and ranging slot assignments 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.

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.

5.1.2 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.

5.1.3 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.

5.1.4 Caveats

The current implementation of global LWB scheduling may not be the optimal choice for large-scale deployments, as the localization update rate for each tag is inversely proportional to the number of tags which are scheduled. In the case of large-scale deployments, tags will likely be able to be separated into disjoint sets of observed anchors, allowing them to occupy the same timeslot without interference. A more thorough discussion of the limitations imposed by our current implementation is described in [Section 8](#).

6. TRIPPOINT MODULE

The TriPoint module, shown in [Figure 6a](#), is a solder-on module that abstracts the complexity of the SurePoint system to provide embedded devices with straightforward access to their location. From a hardware designer's viewpoint, adding TriPoint to a project is as simple as connecting a new integrated circuit (IC). TriPoint requires just power, ground, and three I/O lines, plus three UWB antennas that fit the form factor of the hardware design. From a software engineer's viewpoint, TriPoint is just a new interrupt source that periodically provides location information and can be read over a standard bus. By providing the module abstraction, TriPoint greatly reduces the complexity for system designers to incorporate device position information into their designs.

6.1 TriPoint Module Design

Conceptually, TriPoint is a co-processor and radio that implements the SurePoint system. It includes the DecaWave DW1000 UWB radio [9], a STM32F0 ARM microcontroller [36] for running the software algorithms, an RF switch for changing antennas, and the required regulators, crystals and passives needed by the ICs. All of the SurePoint software runs on the onboard MCU. This includes the broadcast ranging protocol, the flooding-based time synchronization, and continuous online calibration to preserve synchronization.

There are several benefits to this approach. First, each TriPoint module supports both the tag and anchor protocols and can be switched between modes at runtime. This ensures that the entire software suite is on each TriPoint removes the need for a module to be configured or reprogrammed for each design. Second, this approach allows the TriPoint MCU to handle all of SurePoint's timing sensitive operations without limiting or complicating the main application microcontroller's event loop. Third, centralizing the software allows the onboard MCU to mask the complexity and silicon bugs of the DW1000 radio. In our experience, programming the DW1000 is challenging, and providing an abstraction layer above the radio significantly reduces development time.

Physically, the TriPoint module is a triangle shape to facilitate installing antennas at 120° offsets from each other. As mentioned in [Section 2.3](#), mounting the antennas in this orientation helps minimize the probability that all three antennas will cancel due to cross-polarization or antenna nulls.

6.2 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 [8]. 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. SurePoint utilizes a production-time calibration procedure to measure the RX and TX delays of each

Figure 6: **SurePoint hardware.** (a) shows the TriPoint module, a solder-on device that implements the SurePoint system and exposes location information over an I²C 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.

Figure 7: 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 7b.

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.

Command	Description
Info	Read the TriPoint identifier and version number.
Config	Configure tag and anchor options.
Read Interrupt	Query the SurePoint 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 1: I²C interface to TriPoint.

These measurements are sufficient for single-channel two-way time-of-flight calibration. However, for the SurePoint 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.

6.3 TriPoint I²C Interface

The TriPoint module provides an I²C interface plus an interrupt line for configuration and location updates. A set of commands, briefly highlighted in Table 1, are defined as I²C 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 I²C 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.

6.4 TriTag Carrier Board

TriTag, shown in Figure 6b, 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 7a shows how the system can be a smartphone accessory.

7. EVALUATION

We begin our evaluation 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.

7.1 Figure 8: Stationary Tracking

We begin by replicating the cross experiment from PolyPoint in a near-identical evaluation space.¹ In this experiment, we stand for approximately 15 s at 50 positions in a cross in the middle of the evaluation area. For our experiment, we deploy on 9 anchors as opposed to PolyPoint’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 PolyPoint’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 we 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 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.

7.2 Figure 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 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.

7.3 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 11 explores what happens if you eliminate some diversity by reprocessing a data trace and dropping measurements. From the discussion in Section 2.3, one expects a roughly 9 cm median increase in range error, however, these ranges are then fed into a solver, which may be able 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

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

Figure 8: **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. Precision is computed as the 3D Euclidian distance of points from the median x, y, and z coordinates of all samples.

some spatial or temporal case where certain channel/antenna configurations consistently return the best range estimates. Figure 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. [4], 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.

7.4 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 9: **Mobile Tag**. The addition of diversity to the ranging protocol means that capturing a single sample requires 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.

Figure 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 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.

(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 11: Impact of Diversity on Location. To investigate the importance of diversity for localization quality, we take the motion trace from Figure 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.

(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 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.

Figure 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	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 13: Position Solving. Guided by Gezici [16], 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.

7.5 Iterative Solver

As our final improvement over PolyPoint, 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 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.

7.6 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 $T_{round} = 1$ s, $t_{flood} = 10$ ms, and $t_{slot} = 80$ 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 $\left(\frac{11 \times 80 \text{ ms} + 1 \times 120 \text{ ms}}{12} \right)^{-1} = 12$ 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.

Figure 14: **Tag Join Contention.** To evaluate how our system handles multiple tags entering the space concurrently, we attach 10 tags to a powerstrip 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.

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 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.

7.7 Microbenchmarks / TriPoint Eval

As a modular design, the TriPoint module adds some burden to the systems in which it is deployed. First, the TriPoint module is 5 cm² 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 \$1.01 in modest volume.

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 μ A.

8. DISCUSSION

Throughout the previous sections, we have described and evaluated the SurePoint system which has been developed with the purpose of providing reliable, robust, high-fidelity indoor localization estimates with minimal impact on update rate, and multi-tag support. Here we discuss some limitations of the radio along with future room for improvement to better aid in achieving robust operation with increased tag update intervals.

8.1 Transceiver Limitations

There are a number of implementation nuances imposed by the DW1000 IC that affect the performance of SurePoint. First, the DW1000's SPI speed limitation caps SPI transactions at 20 MHz. This similarly caps the speed of continuous packet receptions to around 1 kHz. Given the short packet durations ($\approx 100 \mu$ s), multiple tags could be time segmented to achieve up to a 10 $\times$ improvement in localization update rate.

8.2 Future Work

The current SurePoint protocol assumes equal localization update requirements across all tags in the network. However, there may times in which certain nodes may be under strict power consumption 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.

9. RELATED WORK

SurePoint spans many domains. The first contribution improves UWB ranging performance. Traditional approaches to improving ranging have come from various mechanisms for better estimating the time of arrival of the line of sight path from a single channel measurement [12, 28]. SurePoint's frequency and polarization diversity is complimentary to these efforts, capitalizing on the improved measurements for each of its diversity samples.

As an indoor localization system, SurePoint is in good company. Table 2 presents an overview of recent commercial and research indoor localization technologies. SurePoint is competitive with the state-of-the-art precision, accuracy, and SWaP (size, weight, and power) metrics. SurePoint differentiates itself, however, by adding support for multiple concurrent tags in the same physical space while maintaining a high update rate and support for highly mobile tags.

Broadcast-based localization systems such as Lazik's ALPS [27] and Kuo's Luxapose [24] naturally support an unlimited number of tags, but those systems are limited by the top speed of tags they can track, by motion blur for Kuo's imager or the long (400 ms minimum) sampling window of Lazik's ultrasonic receiver. Broadcast-based systems have the advantage that users maintain privacy, the system need not know they are there, with the complimentary disadvantage that the building is unaware of the people and things in its space and cannot optimize for them.

Constructive interference is a relatively new phenomenon in the sensor networking community. Born out of Backcast [11] and then shown to be viable for systems by Glossy [15], we are only beginning to explore the potential for constructive interference. As the phenomenon relies on the implementation details of the underlying modulation scheme, we are pleased to show that it continues to be viable under BPM-BPSK as used by the 802.15.4 UWB PHY. To the best of our knowledge, SurePoint is the first system to demonstrate ultra-wideband constructive interference.

Our multi-tag protocol is inspired by the Low Power Wireless Bus (LWB) [14]. Unlike LWB, however, SurePoint's principle operation

System	Technology	Precision	Accuracy	Update Rate	Multiple Tags?	Top Tag Speed	Tag Power	Tag Volume	Max Tag/Anchor Dist
WASP [35]	NB (5.8 GHz) ToA	16.3 cm	50 cm (82%ile)	10 Hz	Yes	Several m/s	2-2.5 W	<i>Not Published</i>	<i>Not Published</i>
UbiSense [38]	UWB TDoA+AoA	99% w/in 30 cm	15 cm	33.75 Hz	Yes	<i>Not Published</i>	<i>Not Published</i>	24.5 cm ³	160 m
TimeDomain [2]	UWB TW-ToF	2.3 cm	2.1 cm	150 Hz	Yes	<i>Not Published</i>	4.2 W	97 cm ³	“hundreds of m”
Lazik et. al [27]	Ultrasonic TDoA	<i>Not Published</i>	3 cm (med) 12 cm (90%)	0.9 Hz	Yes	<i>Not Published</i>	1.1 W [¶]	88 cm ³	100 m
Harmonia [19]	UWB TDoA	<i>Not Published</i>	39 cm (med) 82 cm (90%)	56 Hz	No	<i>Not Published</i>	120 mW ^{**}	<i>Not Published</i>	<i>Not Published</i>
Tagoram [42]	NB (UHF) SAR	<i>Not Published</i>	12.3 cm (med)	At most 30 Hz	No	0.5 m/s	N/A	8 cm ³	10 m
WiTrack [3]	UWB ToF	<i>Not Published</i>	12 cm (med) 31 cm (90%)	At most 400 Hz	No	<i>Not Published</i>	N/A	32,700 cm ³ (avg torso [6])	(<i>Not Published</i>) > 11 m
RF-IDraw [40]	NB (UHF) Interferometry	3.6 cm (med) 3.7 cm (90%)	19 cm (med) 38 cm (90%)	At most 53 Hz	No	0.5 m/s [*]	N/A	8 cm ³	9 m
PolyPoint [20]	UWB ToF	31 cm	39 cm (med) 140 cm (90%)	16 Hz	No	<i>Not Published</i>	150 mW	9 cm ³	50 m
Harmonium [21]	UWB TDoA	9 cm (med) 16 cm (90%)	14 cm (med) 31 cm (90%)	19 Hz	No	2.4 m/s ^{††}	75 mW	1.5 cm ³	78 m
Chronos [39]	Bandstitched UWB ToF	<i>Not Published</i>	65 cm (med) 170 cm (90%)	12 Hz	No	<i>Not Published</i>	1.6 W [§]	2.7 cm ^{3¶}	<i>Not Published</i>
SurePoint	UWB ToF	12 cm (med) 28 cm (90%)	29 cm (med) 50 cm (90%)	1-12 Hz	Yes	at least 2.4 m/s	280 mW	3 cm³	50 m

[§] Using reported power numbers from [17] for Intel WiFi Link 5300 in RX mode. [¶] Assuming smaller, PCIe Half Mini Card form factor.

^{¶¶} Estimate from power draw of similar audio+network apps [25]

^{††} Estimated as $(56 \text{ Hz} / 3.5 \text{ GHz} \times c) / 2$

^{§§} The paper reports only “real-time”, however this is as perceived by a human user, which may not be sufficient for applications such as controls.

^{*} This paper reports no speed information, but uses the same tag and similar anchors as Tagoram, so we use the same top speed estimate.

^{**} Published power draw of 8.5 mW is in addition to a traditional narrowband radio. This estimate adds a CC2520 as a representative low-power radio.

Table 2: Comparison of localization quality, utility, and SWaP performance for recent high-performing indoor RF localization systems. Where possible, reasonable extrapolations are made. SurePoint achieves comparable localization performance with best in class systems, exceeding several in through-wall cases, with near-best SWaP metrics, from independent measurements capable of tracking faster-moving objects than nearly any other system.

is ranging, not data sharing. While LWB floods every packet, SurePoint floods only the schedule and contention slots (when new tags join). During ranging, a tag communicates only with its one-hop neighbors. The scheduling slots ensure that there is no interference during ranging. As SurePoint anchors are wall-powered and communication is only meaningful in areas that are well-covered (otherwise there are insufficient anchors to range), tags need not actually participate fully in floods, they simply need to send or hear one packet and then may drop out, saving tag battery life without meaningfully detracting from the flood.

10. CONCLUSIONS

We present SurePoint, a new indoor localization system that achieves decimeter-level accuracy and maintains sub-meter 99th percentile accuracy. We demonstrate how the addition of simple spatial, frequency, and polarization diversity improves the quality and reliability of ultra-wideband range estimates. We then showcase an efficient and robust broadcast ranging protocol for collecting high-fidelity ranging estimates from numerous tags. We then demonstrate that the constructive interference phenomenon remains effective for the 802.15.4 UWB channel, and leverage this to support localization of multiple tags. Finally, we introduce TriPoint, a hardware module that encapsulates all the complexities of ranging, reducing localization to a simple request.

11. ACKNOWLEDGMENTS

This research was conducted with Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

12. REFERENCES

- [1] Nanotron RTLS. http://nanotron.com/EN/PR_find.php.
- [2] Time Domain PulsON 400 RCM. <http://www.timedomain.com/p400.php> and http://www.timedomain.com/datasheets/TD_DS_P410_RCM_FA.pdf.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] C. E. Clauser et al. Weight, volume, and center of mass of segments of the human body. Technical report, Air Force Systems Command, 1969.
- [7] 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.
- [8] DecaWave. Antenna delay calibration of DW1000-based products and systems.
- [9] DecaWave. ScenSor DW1000. <http://www.decawave.com/>.
- [10] E. Dhalgren and H. Mahmood. Evaluation of indoor positioning based on Bluetooth Smart technology. Master’s thesis, Chalmers, Jun 2014.

- [11] 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.
- [12] 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*, (1):1–13, 2006.
- [13] Federal Communications Commission. First report and order 02-48. Technical report, February 2002.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall. Demystifying 802.11n power consumption. In *Proceedings of the 2010 International Conference on Power Aware Computing and Systems*, HotPower’10, pages 1–, 2010.
- [18] IEEE Computer Society. Part 15.4: Low-rate wireless personal area networks (LR-WPANs). Technical report.
- [19] 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.
- [20] 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.
- [21] 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, April 2016.
- [22] K. Klipp, J. Willaredt, H. Rosé, and I. Radusch. Low cost high precision indoor localization system fusing inertial and magnetic field sensor data with radio beacons. Second Annual Microsoft Indoor Localization Competition, April 2015.
- [23] 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.
- [24] 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.
- [25] Y. M. Kwete. Power consumption testing for iOS. Master’s thesis, North Dakota, October 2013.
- [26] P. Lazik, N. Rajagopal, O. Shih, B. Sinopoli, and A. Rowe. ALPS: A bluetooth and ultrasound platform for mapping and localization. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, SenSys ’15, pages 73–84, New York, NY, USA, 2015. ACM.
- [27] 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.
- [28] 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.
- [29] D. Lymberopoulos, J. Liu, X. Yang, A. Naguib, A. Rowe, N. Trigoni, and N. Moayeri. Microsoft indoor localization competition – IPSN 2015.
- [30] D. Lymberopoulos, J. Liu, Y. Zhang, P. Dutta, Y. Xue, and A. Rowe. Microsoft indoor localization competition – IPSN 2016.
- [31] 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. Technical report, 2004.
- [32] S. Patel, K. Truong, and G. Abowd. PowerLine Positioning: A practical sub-room-level indoor location system for domestic use. In *Ubiquitous Computing*, UbiComp’06, pages 441–458. Springer Berlin Heidelberg, 2006.
- [33] 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.
- [34] C. Sánchez, S. Ceriani, P. Taddei, E. Wolfart, and V. Sequeira. STEAM sensor tracking and mapping. Second Annual Microsoft Indoor Localization Competition, April 2015.
- [35] 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.
- [36] STMicroelectronics. STM32F031G6U6. <http://www.st.com/web/catalog/mmc/FM141/SC1169/SS1574/LN7/PF259775>.
- [37] 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.
- [38] Ubisense. Series 7000 Compact Tag. http://www.ubisense.net/en/media/pdfs/products_pdf/uk/80553_series_7000_compact_tag.pdf.
- [39] 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.
- [40] 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.
- [41] J. Wright and S. Nocedal. *Numerical Optimization*. SciPy, 1999.
- [42] 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.