

PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization

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

ABSTRACT

We introduce PolyPoint, the first RF localization system which enables the real-time tracking and navigating of quadrotors through complex indoor environments. PolyPoint leverages the new ScenSor transceiver from DecaWave to acquire the timestamps necessary for accurate time-based location estimation and leverages the benefits of antenna and frequency diversity to iteratively refine a tag's position. PolyPoint produces quadrotor position estimates at a rate of 20 Hz with median error below 40 cm and average error of 56 cm in line-of-sight conditions. PolyPoint approaches the localization accuracy necessary to safely navigate quadrotors indoors, a feat currently achieved by costly and delicate optical motion capture systems.

1. INTRODUCTION

Global Navigation Satellite Systems (GNSS) have long been the universal gold standard for the accurate navigation of outdoor spaces. Certain systems have enabled position determination accuracy better than 10 cm, enabling the unassisted navigation of quadrotors outdoors—an emerging area of interest with applications ranging from aerial photography to express courier services. Unfortunately, these RF systems break down in indoor environments due to extreme attenuation and heavy multipath, necessitating the use of local navigation aids. Recently-available commodity ultra-wideband RF transceivers show promise in supporting the ranging accuracy required for quadrotor navigation, but they still suffer from large ranging errors if the RF line-of-sight path is attenuated. Attenuation of the line-of-sight path can occur when there exists a polarization mismatch between antennas. The simple addition of antenna and channel diversity can allow for a significant improvement in average ranging accuracy indoors, as much as 27 cm to 3 cm, with little increase in overall system cost or complexity. This simple addition elevates RF localization system performance to the level necessary to enable quadrotor navigation indoors—a feat previously achieved only by fragile and costly optical motion capture systems [10, 14].

Optical motion capture systems track reflective markers [10, 14]. Cameras mounted in known locations around the environment track the position of any reflective surfaces within their field of view and are able to determine a quadrotor's position if an affixed reflective

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

HotWireless '15, Sept 11, 2015, Paris, France
ACM 978-1-4503-3699-4/15/09.
<http://dx.doi.org/10.1145/2799650.2799651>

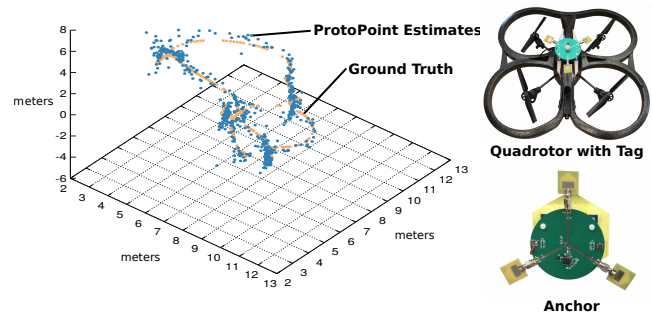



Figure 1: PolyPoint reconstructs the flight path of a Parrot AR.Drone in a $15 \times 15 \times 15$ m space at a 16 Hz update rate using sixteen identical nodes—fifteen fixed anchors and one mobile tag. Ground truth is acquired using the commercial OptiTrack [10] optical motion capture system. PolyPoint tracks the path with 39 cm median error and 140 cm 95th percentile error.


marker is seen by two or more cameras at once. With sufficiently high-resolution cameras observing a reasonable volume, optical systems can achieve localization accuracy better than 1 cm. They can be sensitive to lighting conditions and other reflective surfaces situated throughout the environment, however. This along with the high cost (typically \$1,000–\$10,000 per camera) of these systems has hindered their use beyond the studio and laboratory.

In contrast, RF-based techniques for indoor localization offer clear advantages in cost and reliability, but systems built on commodity narrowband radios have not yet shown the level of accuracy required for unassisted quadrotor flight or require movement to improve localization accuracy through post-processing tasks [15]. The low utilized bandwidth hinders the ability of these systems to distinguish the effects of the line-of-sight path from the effects of multipath, thus significantly impacting their performance in heavy multipath environments. By increasing the utilized bandwidth, the multipath resolution improves, leading to the use of ultra-wideband signals in systems that require high accuracy in indoor environments [8].

Ultra-wideband location systems such as those offered by Time-Domain [2] and Ubisense [13] have long held an important role in providing accurate real-time location services for defense, robotics, and industrial applications. However, the high system cost has prohibited their widespread adoption outside these niche markets. This has changed, however, with the recent release of the DecaWave ScenSor chipset [5], which integrates an ultra-wideband transceiver and targets indoor localization applications. The transceiver provides packet receive and transmit timestamps with up to 16 ps of precision, primitives that enable a variety of custom-tailored localization



(a) Node Architecture



(b) Node Implementation

Figure 2: PolyPoint node architecture and implementation. PolyPoint nodes consist of a DecaWave DW1000 ScenSor transceiver for data transmission and timestamping, a SKY13317 RF switch and three UWB antennas for diversity, and a CC2538 system-on-chip [12] for protocol orchestration and data offload. PolyPoint uses the same hardware for both the tag and anchors, leading to a simple and symmetric design intended to aid further research in lower-level details such as protocol and hardware design for use in tailored applications.

protocols. Furthermore, the transceivers are inexpensive, allowing for the possibility of ubiquitous indoor localization.

Ultra-wideband transceivers such as ScenSor still require the line-of-sight path to be unobstructed to produce accurate time-of-arrival measurements. In real-world scenarios, however, there are a variety of factors which can cause the line-of-sight path to be obstructed. While sufficient node density and careful placement can help avoid obstructions, differences in polarization between nodes can lead to significant attenuation of the line-of-sight path without the presence of any interfering objects. For this reason, we introduce multiple antennas at each node to mitigate the effects of polarization mismatch on ranging (and subsequently localization) error.

In this paper, we introduce PolyPoint, a new localization system that couples the DecaWave ScenSor transceiver with antenna diversity and a new, efficient ranging protocol. By utilizing antenna diversity, PolyPoint shows an order of magnitude improvement in accuracy over the use of just one antenna at each node. Furthermore, PolyPoint’s custom ranging protocol supports the aggregation of many different range estimates across different antenna and RF channel combinations while still maintaining an update rate at tens of Hz, necessary for tracking fast-moving objects. The following sections detail the design, implementation, and evaluation of the PolyPoint system and its performance when tracking quadrotors indoors.

2. OVERVIEW

PolyPoint is a localization system that uses one-way time-of-flight measurements to derive range estimates between a mobile tag and fixed-location infrastructure (anchors) to determine a tag’s position. By collecting range estimates between the tag and three or more anchors, the system is able to calculate the tag’s 3D position using trilateration. The mobile tag is affixed to the object to track, and position estimates are calculated in real-time.

2.1 PolyPoint Hardware Design

The PolyPoint hardware design, shown in Figure 2, targets a symmetric and modular architecture for ease in system development and evaluation. Connections between the controller and radio as well as the connections to each antenna are interchangeable to enable the use of differing controller and antenna designs without a change in PolyPoint’s core hardware architecture.

The PolyPoint RF hardware consists of the ScenSor IC for data communication and ranging operations along with an RF switch and three UWB antennas¹ for antenna diversity. The three antennas are oriented at 0, 120, and 240 degrees to enable adequate coverage of polarization. An Atum sensor node [1] along with its associated breakout board is used to orchestrate ranging operations and offload range estimates over USB or 802.15.4. Finally, the system is powered via USB connection to a PC or wall charger.

2.2 PolyPoint Software Design

All ranging operations are orchestrated using an Atum sensor node which is built around the CC2538 system-on-chip. The embedded ARM core is used to run the Contiki embedded operating system and performs timing-critical operations for the ScenSor chipset [6]. The tag node periodically requests range estimates from all anchors within range. Once range estimates to all nearby anchors have been collected, the range estimates are offloaded to a connected PC to calculate the tag’s position.

To leverage the advantages of antenna and channel diversity, PolyPoint combines 27 successive ranging estimates. One ranging estimate is obtained for each combination of three tag antennas, three anchor antennas, and three different RF channels. Figure 3 shows the distribution of ranging error across a large number of range estimates. Empirically, we find that the 10th percentile of ranges estimates gives the least average error. Therefore, the 10th percentile aggregate of the 27 range estimates is used as the true range at each anchor.

Deriving a high-precision time-of-flight range estimate between two unsynchronized nodes requires *three* packets per range pair. Basic unsynchronized time-of-flight only sends a POLL from a tag to an anchor and a RESPONSE from an anchor to a tag:

$$\begin{aligned} & \textcircled{1} \text{ Tag} \xrightarrow{\text{Poll}} \text{Anc} \quad \textcircled{2} \text{ Anc} \xrightarrow{\text{Resp}} \text{Tag} \\ T_oF = & [(\text{Tag}_{RX_Resp} - \text{Tag}_{TX_Poll}) - \\ & (\text{Anc}_{TX_Resp} - \text{Anc}_{RX_Poll})] / 2 \end{aligned}$$

Critically, however, this assumes that the clocks on the tag and anchor are running at exactly the same frequency. To compensate for crystal variances, one must send an additional packet, REF, usually before sending POLL. Adding in crystal correction, unsynchronized

¹ UWB antenna design borrowed from Azim et. al [3].

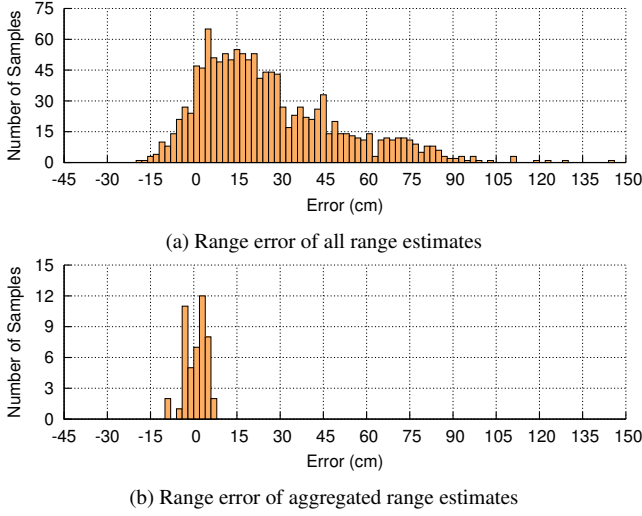


Figure 3: The distribution of range estimation error before and after aggregation. The tag collects 27 range estimates to each anchor, across the 27 combinations of tag antenna, anchor antenna, and RF channel. Before aggregation, range error is large and predominantly positive. This informs the selection of the 10th percentile of the 27 ranges as an aggregation mechanism. After aggregation, range error is significantly lower, leading to an improvement in the overall localization accuracy.

time-of-flight is now

$$\begin{aligned}
 & \textcircled{1} \text{ Tag} \xrightarrow{Ref} \text{Anc} \quad \textcircled{2} \text{ Tag} \xrightarrow{Poll} \text{Anc} \quad \textcircled{3} \text{ Anc} \xrightarrow{Resp} \text{Tag} \\
 K &= \frac{\text{Tag}_{TX_Poll} - \text{Tag}_{TX_Ref}}{\text{Anc}_{RX_Poll} - \text{Anc}_{RX_Ref}} \\
 ToF &= [(\text{Tag}_{RX_Resp} - \text{Tag}_{TX_Poll}) - \\
 & K * (\text{Anc}_{TX_Resp} - \text{Anc}_{RX_Poll})] / 2
 \end{aligned}$$

In a system with the minimum 3 anchors, a naïve protocol requires $27 \times 3 \times 3 = 243$ packets. In order to reduce the total number of message exchanges, PolyPoint tags transmit a broadcast for each of the 27 different configurations. The difference in range estimates can be calculated at each anchor by observing the time differences between successive message receptions. A final two-way time-of-flight exchange is then performed between the tag and each anchor to account for any error in range due to clock offset. Figure 4 shows the modified two-way time-of-flight protocol used to quickly gather each of the 27 successive range estimates. For this revised protocol, 3 anchors require only $27 + 1 + 3 = 31$ packets.

To minimize transmission overhead, tags locally compute the 10th percentile range to each anchor, and then transmit these aggregate range estimates over 802.15.4 to a nearby PC for processing. PolyPoint uses a non-linear least squares approach to combine the range estimates of all nearby anchors [9]. From here, the position estimates can either be stored for offline analysis or used as feedback to control the position of the object under observation.

3. EVALUATION

We design two experiments to evaluate the accuracy and precision of PolyPoint in an indoor environment and its suitability for the intended application of tracking quadrotors indoors. The tests are run in a large 20×20 m indoor space. Ground truth is captured with

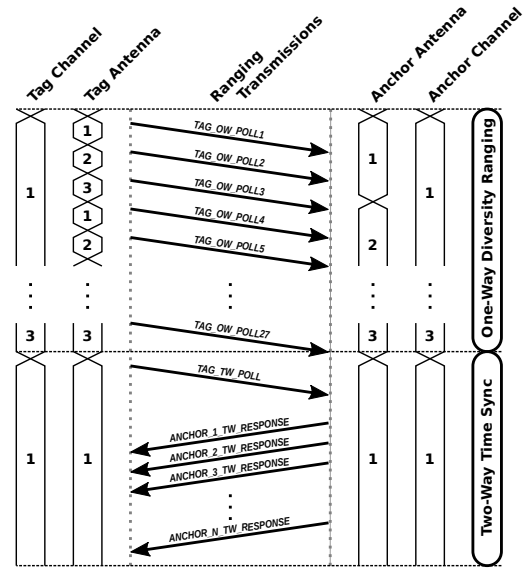


Figure 4: Ranging protocol. PolyPoint is able to leverage antenna diversity without significantly impacting position update rate. The protocol starts with a series of 27 broadcast transmissions from the tag for each combination of tag antenna, anchor antenna, and RF channel. The time-of-arrival data 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.

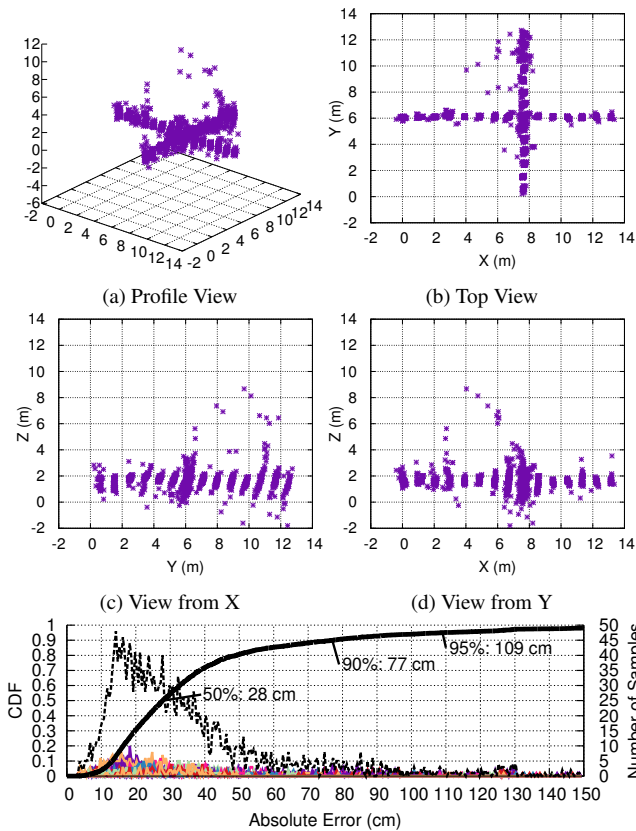
an OptiTrack motion capture system [10], calibrated to a reported accuracy of better than 1 mm across the evaluation space.

To evaluate the accuracy and precision of position estimates offered by PolyPoint, along with an estimate of PolyPoint’s performance when applied to the localization of stationary or slowly-moving objects, we measure the distribution of position estimates across a variety of known stationary locations. Figure 5 shows the point cloud of position estimates at 29 known locations arranged in the shape of a cross spanning the evaluation space. Over the points covered by this test, PolyPoint achieves an average accuracy of 28 cm with a precision of 31 cm.

The second test evaluates PolyPoint’s performance at tracking the real-time position of a quadrotor. For this test, we affix a PolyPoint node to a Parrot AR.Drone 2.0 [11] quadrotor and manually fly it through the indoor environment. In offline analysis, we compare the reconstructed path provided by PolyPoint to the ground truth observations to estimate average localization error. The flight path and results are captured in Figure 6. PolyPoint achieves an average error of 56 cm throughout the flight test.

4. DISCUSSION AND FUTURE WORK

By leveraging the benefits of UWB signals and antenna diversity, we show that PolyPoint provides best-in-class indoor localization accuracy for RF localization systems built on commodity radios. There are, however, several aspects of PolyPoint that warrant further research and consideration.



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

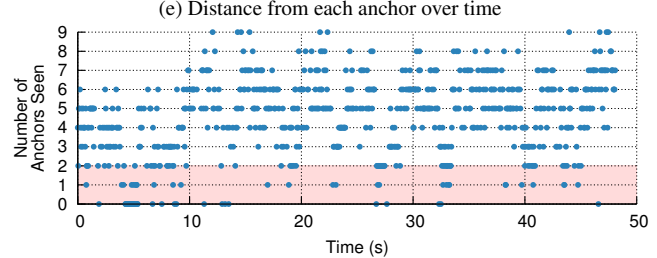
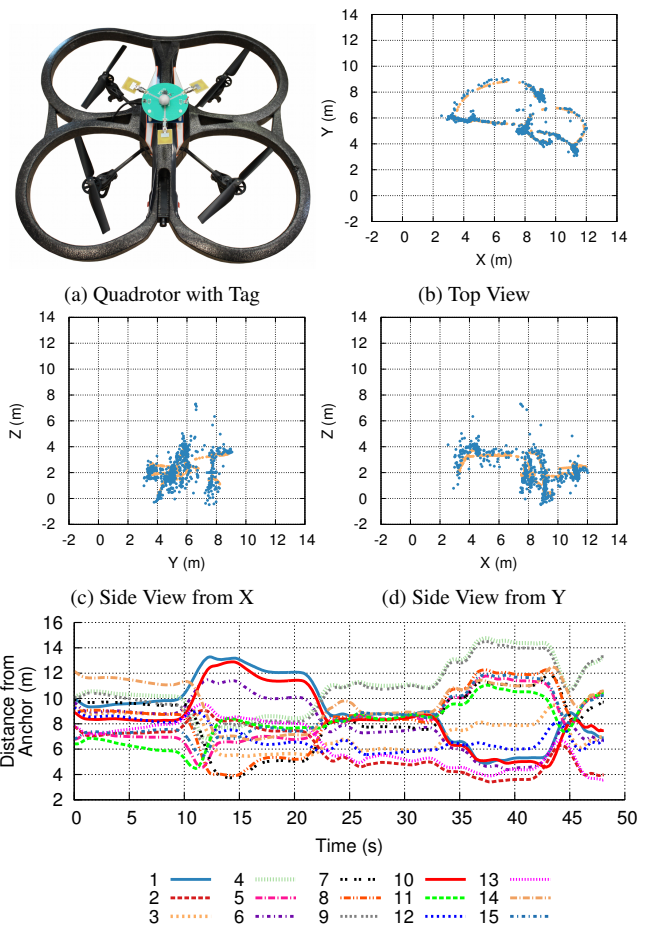
Figure 5: Stationary tracking experiment

4.1 Line-of-Sight Limitations


The indoor space and anchor placement selected for evaluation in this paper were chosen to avoid the presence of occluding obstacles between the tag and each anchor. The presence of more clutter within the environment or complex environment geometries can lead to instances of non line-of-sight (NLoS) propagation. Instances such as these lead to greater than expected range estimates because a later (reflecting) path's time will inadvertently be used instead of the desired direct path. UWB only enables PolyPoint to distinguish between the LoS and NLoS path when both are present. A variety of methods have been proposed in the past to deal with instances of NLoS [4, 7], and similar methods could be employed to aid in the identification and removal of the affected range estimates.

4.2 Software Limitations


The ranging protocol outlined in this paper has been tailored to aggregate numerous observations of range across varying antenna and RF channel arrangements with minimal impact on the system's update rate. There are, however, various details in the protocol's implementation which could be improved to positively affect overall system performance. First, the current protocol implementation relies on the reception of the first one-way range estimate in order to successfully complete a ranging request. Furthermore, the two-way time-of-flight messages must all be received without error in order to complete a single ranging operation. Since all of these messages are communicated using the same combination of tag antenna, anchor antenna, and RF channel, a significant performance penalty can be incurred if this particular channel experiences significant attenuation. This limitation can cause a reduction in the number of anchors re-




(e) Distance from each anchor over time



(f) Number of anchors seen over time



(g) Error over time



(h) Error CDF and Histogram

Figure 6: Quadrotor tracking experiment




Figure 7: CDF of packet inter-arrival time. The tag reports measurements over a lossy 802.15.4 link. The tag sends a result packet every round, even if zero anchors are seen. Round duration can vary slightly depending on the number of anchors that reply. The steps correspond to increasing numbers of sequentially dropped packets. Extrapolating from the average arrival rate when packets are not dropped, we would expect to receive about 990 packets, however only 714 were received, a packet reception rate of only 72%.

porting range, leading to an increase in position error. In the future, PolyPoint could be modified to provide redundant transmissions with varying antenna selections or provide additional mechanisms to choose the best antenna selection on-the-fly.

One major limitation of the current PolyPoint protocol lies in its fixed time-slotting assignments. This methodology does not scale well for systems with more than a handful of nodes. While numerous neighbor discovery protocols exist, few handle highly mobile and transient nodes, such as a quadrotor traversing a building. One approach could run traditional neighbor discovery on the static anchor network and use successive tag position estimates to proactively predict the neighbors that will arrive and disappear.

4.3 Hardware Limitations

The current hardware implementation for PolyPoint is bulky, which limits its use on anything but the least constrained quadrotor devices. The payload mass contributes directly to the active power consumption, and therefore has the potential to greatly impact the overall quadrotor flight time. The mass of the current PolyPoint node is 95 g. A majority of this mass can be attributed to discrete cabling, RF and digital connectors, and the microcontroller’s breakout board. By shrinking the overall design to fit within a manageable 5×5 cm² area and using direct-to-board RF connections, the system mass could be reduced to around 15 g. This would greatly open up the possibility of accurate and responsive quadrotor localization on a more aggressively constrained class of ‘micro-quadrotors’.

4.4 Position Solving and Communication

Currently, the tag offloads range estimates and leverages a local cloudlet to solve for the final position. This approach incurs a modest additional communication latency, but also introduces a surprisingly large new source of error. We configure the PolyPoint tag to send each new position message in a 802.15.4 broadcast messages, which are collected by a powered base station never more than 15 m from tag. Figure 7 shows the packet inter-arrival time for the quadrotor trace. Despite the fairly ideal conditions, the 802.15.4 channel exhibited over 20% packet loss, resulting in potentially many missed position estimates. A reliable transport layer would greatly improve the positioning robustness.

5. CONCLUSIONS

PolyPoint is the first RF localization system which has shown the accuracy and responsiveness required for the suitable navigation of quadrotors indoors. PolyPoint leverages the accurate ranging capa-

bilities afforded by the first commodity ultra-wideband transceiver along with the simple addition of antenna diversity to produce accurate, real-time estimates of location indoors. Additionally, PolyPoint has introduced a new ranging protocol which minimizes the number of message transmissions required to leverage the additional antenna diversity at each node.

The advent of decimeter-scale indoor localization accuracy opens up the door to many new areas of research and commerce. Innovations such as hybrid outdoor and indoor localization approaches may help to aid in the realization of a global unified localization platform. Other areas of research such as self-localization and self-discovery protocols are now feasible and will be necessary to help aid in the deployment of these systems at scale in the future.

6. 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.

7. REFERENCES

- [1] Atum Sensor Node. <https://github.com/lab11/atum>.
- [2] Time Domain PulsON 410 RCM. <http://www.timedomain.com/p400.php>.
- [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] J. Borras, P. Hatrack, and N. B. Mandayam. Decision theoretic framework for NLOS identification. In *Vehicular Technology Conference, 1998. VTC 98. 48th IEEE*, volume 2, pages 1583–1587. IEEE, 1998.
- [5] DecaWave. ScenSor DW1000. <http://www.decawave.com/>.
- [6] A. Dunkels, B. Gronvall, and T. Voigt. Contiki – a lightweight and flexible operating system for tiny networked sensors. In *Local Computer Networks, 2004. 29th Annual IEEE International Conference on*, pages 455–462, Nov 2004.
- [7] I. Guvenc, C.-C. Chong, and F. Watanabe. NLOS identification and mitigation for uwb localization systems. In *Wireless Communications and Networking Conference, 2007. WCNC 2007. IEEE*, pages 1571–1576. IEEE, 2007.
- [8] 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.
- [9] W. Murphy and W. Hereman. Determination of a position in three dimensions using trilateration and approximate distances. *Department of Mathematical and Computer Sciences, Colorado School of Mines, Golden, Colorado, MCS-95*, 7:19, 1995.
- [10] NaturalPoint. OptiTrack. <http://www.optitrack.com>.
- [11] Parrot. AR.Drone 2.0. <http://ardrone2.parrot.com/>.
- [12] Texas Instruments. CC2538 System-on-Chip. <http://www.ti.com/product/cc2538>.
- [13] Ubisense. Series 7000 Compact Tag. http://www.ubisense.net/en/media/pdfs/products_pdf/uk/80553_series_7000_compact_tag.pdf.
- [14] VICON Motion Capture. <http://www.vicon.com>.
- [15] 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*, pages 235–246. ACM, 2014.