

Radio Interferometric Tracking of Mobile Wireless Nodes

Branislav Kusy
Janos Sallai
Gyorgy Balogh
Akos Ledeczki
Vanderbilt University, USA
akos@isis.vanderbilt.edu

Vladimir Protopopescu
Johnny Tolliver
Frank DeNap
Oak Ridge National Lab, USA
protopopesva@ornl.gov

Morey Parang
University of Tennessee, USA
parangm@ornl.gov

ABSTRACT

Location-awareness is an important requirement for many mobile wireless applications today. When GPS is not applicable because of the required precision and/or the resource constraints on the hardware platform, radio interferometric ranging may offer an alternative. In this paper, we present a technique that enables the precise tracking of multiple wireless nodes simultaneously. It relies on multiple infrastructure nodes deployed at known locations measuring the position of tracked mobile nodes using radio interferometry. In addition to location information, the approach also provides node velocity estimates by measuring the Doppler shift of the interference signal. The performance of the technique is evaluated using a prototype implementation on mote-class wireless sensor nodes. Finally, a possible application scenario of dirty bomb detection in a football stadium is briefly described.

Categories and Subject Descriptors: C.2.4[Computer-Communications Networks]:Distributed Systems

General Terms: Algorithms, Experimentation, Theory

Keywords: Wireless Sensor Networks, Radio Interferometry, Tracking, Localization, Location-Awareness, Mobility

Acknowledgments: This work was supported in part by TRUST (The Team for Research in Ubiquitous Secure Technology, NSF award number CCF-0424422), ARO MURI grant SA5212-11087 and a Vanderbilt Discovery Grant. We would like to express our gratitude to Andras Nadas, Miklos Maroti and Peter Volgyesi for their valuable contributions to this work. We are indebted to Jay Lowe for providing access to the Vanderbilt Football Stadium, the anonymous reviewers and our shepherd, Venkat Padmanabhan, for their constructive comments.

1. INTRODUCTION

The Global Positioning System (GPS) has made the transition from a few early adopters to a sizeable percentage of the population in many technologically advanced countries around the world. Most new cars are offered with built-in GPS systems. Top-of-the-line cell phones have built-in GPS

receivers. On the commercial side, trucking companies have adopted GPS for asset tracking, for example. And of course, the military is using it everywhere from soldier navigation to smart munition guidance. In fact, the European Union announced that they will build a competing system called Galileo to reduce their dependence on American technology.

The advantages of GPS are numerous. It provides accurate positioning anywhere on Earth. Receivers are relatively small, cheap, and power conscious. The service itself is free. On the other hand, GPS has shortcomings too, making it less applicable to certain application domains. The accuracy of traditional GPS is a few meters. Differential GPS can get two orders of magnitude more precise, but it relies on expensive hardware. Typical handheld GPS receivers last only a few hours on a single charge. Recently, single chip GPS receivers became available, but they still cost tens of dollars and typically consume 50-100mW. Moreover, GPS does not work well indoors, in cluttered urban environments or under dense foliage. Finally, GPS can be jammed making it vulnerable in military applications.

Hence, when a wireless mobile application requires geolocation information with one or more of the following properties: (1) high precision (a few centimeters), (2) low per node cost (a few dollars), (3) low power consumption (10-50mW), and (4) robustness to environmental conditions, then alternative solutions need to be applied. We have recently introduced a novel localization method based on radio interferometry [19, 13] primarily targeted at wireless sensor network node localization for static deployments. Later we applied the same underlying technique to tracking a single mobile node [12]. Here we extend this approach to multiple tracked objects and to estimate the velocity, along with the locations of the tracked objects. We also describe a dirty bomb detection and localization system that utilizes the technique.

The underlying idea of radio interferometric ranging is to have two nodes transmit unmodulated radio waves at close frequencies, thereby creating an interference field. Measuring the phase offsets at different points provides information on the location of those points and that of the transmitters. If we have a set of static nodes at unknown positions and we make enough measurements, the relative coordinates of the nodes can be reconstructed. For tracking, we can place several nodes at known positions to create reference points similar in functionality to GPS satellites. The system then uses interferometric ranging to determine the unknown positions of mobile nodes.

Radio interferometric localization and tracking satisfies three of the requirements above: high precision, low cost,

Copyright 2007 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the U.S. Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. *MobiSys'07*, June 11-14, 2007, San Juan, Puerto Rico, USA. Copyright 2007 ACM 978-1-59593-614-1/07/0006 ...\$5.00.

and low power. It can achieve centimeter scale accuracy. Currently, it is implemented on COTS Crossbow Mica2 mote devices with no additional hardware. We utilize the mote’s Texas Instruments CC1000 radio which costs about \$3. This radio could be added to other hardware platforms with an estimated incremental cost of less than \$10. Its receive power requirement is about 20 mW.

The disadvantage of our technique, though, is that it is susceptible to RF multipath, so currently it does not work well indoors. Even though we expect the interferometric ranging to provide accurate localization in large indoors facilities, such as sports arenas where GPS would inevitably fail, at this moment, we do not have an effective localization algorithm for the environments with significant multipath.

Our initial research, however, shows promise in resolving the multipath problem. If multipath is present, the phase offset depends not only on the four distances between the nodes, but also on the relative amplitudes of the reflected signals compared to the direct signal, as well as the locations of the points of reflection. We have mathematically modeled the expected phase offset of the interferometric signal and made the following observation: if we have a bound on the number of reflected signal components and have enough carrier frequencies at our disposal, we will eventually be able to collect more data than the number of unknowns in our theoretical model. This is because with each frequency, we can accurately measure two values: the relative phase offset and the relative signal strength between the two receivers. Note that the signal strength does contain information about the reflected paths. If we measure at slightly different frequencies and hence, wavelengths, the composite signal, consisting of the line of sight component and multiple reflected ones, will have different amplitudes at different frequencies because of the different phases in the summation. As the number of unknowns in our model does not increase, the system will become solvable if enough measurements are available.

Another advantage of our technique is that it is relatively easy to experiment with multiple antenna designs, different radio frequencies, and densities of the infrastructure nodes. Furthermore, the potential of using radio interferometry with lower frequencies which are less susceptible to reflections due to their longer wavelength may significantly simplify the algorithms for indoor localization. To summarize, we believe that radio interferometric ranging is an important technique showing great promise to overcome the main GPS drawback – the lack of indoor localization. Note, however, that we do not claim that our system is a replacement of GPS, as we can only cover a geographically limited area.

The rest of the paper is organized as follows. First we review related work and briefly summarize the theory behind interferometric ranging and highlight previous results. Then we present the description, analysis and evaluation of the enhanced technology for tracking multiple moving nodes simultaneously. We also describe application of our tracking techniques to a dirty bomb detection system and its demonstration in a football stadium. Finally, we conclude the paper with a discussion and future research directions.

2. RELATED WORK

Object tracking systems are used to determine the trajectory of one or more moving targets from partial location information provided by sensors. In general, the location

of an object in the physical world is inferred from different sensed signals, such as acoustic, infrared, or seismic signals. Algorithms for object detection, classification and identification are, therefore, important part of tracking systems. For example, Acoustic ENSBox [7] was used to detect and localize different animal species such as marmots or dusky antbirds. However, the majority of the wireless sensor network (WSN) based tracking systems eliminate the need of object detection by assuming that the tracked nodes cooperate with the system [8, 23, 9, 1].

Cooperative tracking systems need to address multiple issues. First, the systems need to decide which set of sensors to activate to track a particular target and where to send the collected location data. This problem can be solved in an information centric [28] or location centric [3] way. Next, the activated sensors collect ranging measurements to estimate the location of the target, utilizing acoustic signals [7, 25], radio signal strength [9, 1, 27], radio time of flight [6, 16], or phase difference [19]. Finally, filtering and smoothing techniques, such as extended Kalman filters [2, 22] or particle filters [26, 11], are applied to the calculated location estimates to remove the effects of the measurement noise and improve tracking accuracy.

In our work, we concentrate solely on the localization and ranging problem of the tracking systems. In particular, we focus on obtaining highly accurate location estimates that can be calculated fast, to allow for high location update rates. In this regard, the best published results were achieved with systems based on acoustic ranging techniques. Acoustic ENSBox [7] achieved 5 cm average 2D localization error in a large scale experiment, with 1.5 degree average orientation error. Cricket nodes were used to achieve a few cm 2D localization error with 1 Hz update rate at short ranges [21]. SLAT achieved 7 cm 3D localization error indoors [25]. While the achieved localization accuracy is impressive, these systems require hardware components which have considerable cost, form factor, and power consumption, if they are required to work over large distances. State-of-the-art GPS receivers are facing similar problems, despite their relatively low price (\$50) and power consumption (50–100 mW). Due to their small hardware cost overhead, the RF based systems that use the radio chips included on the sensor nodes are an attractive option for tracking systems. Unfortunately, radio signal strength systems do not show a great promise in achieving high location accuracy [27]. A promising sensor network RF time of flight solution described in [16] proposed to utilize pairwise round trip time of flight measurements to relax the strict global time-synchronization requirement and achieved an average error of 0.9 m outdoors and 2 m indoors. Despite the significant progress in many of these ranging approaches, we believe that the radio interferometric ranging provides a superior alternative due to its low-cost (requires only CC1000 radio chip), low power (20 mW receiving power), and high accuracy (4 cm) over large ranges (100 m) [13].

The main problem of tracking moving objects is that the target changes its location during the measurement as the ranging measurement takes a finite amount of time. One solution is to divide the global space-time region into space-time cells [17]. The size of these cells should approximate a region where the signature of the target remains nearly constant. Sequential Monte Carlo localization was shown to achieve improved accuracy of localization when utilizing the

movement of sensor nodes [11]. However, to the best of our knowledge, none of the methods utilized Doppler frequency shifts to calculate the velocity vectors of the target node and improve the localization accuracy simultaneously.

3. TRACKING A SINGLE MOBILE NODE

Radio-interferometry is a novel localization method that measures the relative phase offset of interfering radio signals. Due to the hardware constraints in wireless sensor networks, it is impossible to measure the phase difference of high frequency radio signals with sufficient accuracy. In [19], we describe the Radio-Interferometric Positioning System (RIPS) that measures the phase difference indirectly. This was achieved by analyzing the phase of the low frequency beat signal generated by two interfering radio transmissions and measured using the Radio Signal Strength Indicator (RSSI) signal of the RF transceiver. Highly accurate localization results over large distances (twice the radio communication range) were reported using COTS hardware (XSM by Xbow [4]).

Interferometric ranging was successfully applied in the design of the inTrack tracking system [12]. inTrack reported real-time tracking of a single object with less than one meter localization error using 12 anchor nodes in an 80×90 m area.

Next, let us summarize the theoretical foundations of interferometric ranging and analyze the main weaknesses of the inTrack system that motivated the work described in this paper.

3.1 Interferometric Ranging

Interferometric ranging requires two nodes, A and B , to transmit sine waves at two close frequencies f_a, f_b at the same time. The interference of these two signals results in a beat signal with a beat frequency of $|f_a - f_b|$. The phase difference of the beat signal at two receivers C and D is a function of the distances between the four nodes A, B, C, D :

$$\Delta\phi = \frac{2\pi}{\lambda}(d_{AD} - d_{BD} + d_{BC} - d_{AC})(\text{mod } 2\pi) \quad (1)$$

where $\Delta\phi$ is the difference in phases measured at C and D , $\lambda = \frac{2c}{f_A + f_B}$ is the wavelength corresponding to the mean carrier frequency of the interference signal, and d_{XY} is Euclidean distance of the nodes X and Y [19] (see Fig. 1).

The unique characteristic of RIPS is that it measures so called q -ranges, defined as

$$d_{ABCD} = d_{AD} - d_{BD} + d_{BC} - d_{AC}. \quad (2)$$

An easy to see rearrangement of equation (1) states that the phase difference measured with respect to the wavelength gives us the information about the q -range, i.e. $\Delta\phi = \frac{\Delta\phi}{2\pi}\lambda = d_{ABCD}(\text{mod } \lambda)$. Note that several measurements at different transmit frequencies f_A and f_B are required to resolve the $(\text{mod } \lambda)$ -related ambiguity of the q -range. The q -range can then be found as the solution of the following system of $n \in N$ equations, for $i = 0, \dots, n - 1$:

$$\Delta\phi_i = d_{ABCD}(\text{mod } \lambda_i). \quad (3)$$

The two major difficulties of implementing RIPS on the low-cost, general purpose hardware (such as the XSM mote) are: (1) to set up the radio transceivers to transmit at precise frequencies and (2) to measure the phase of the beat signal

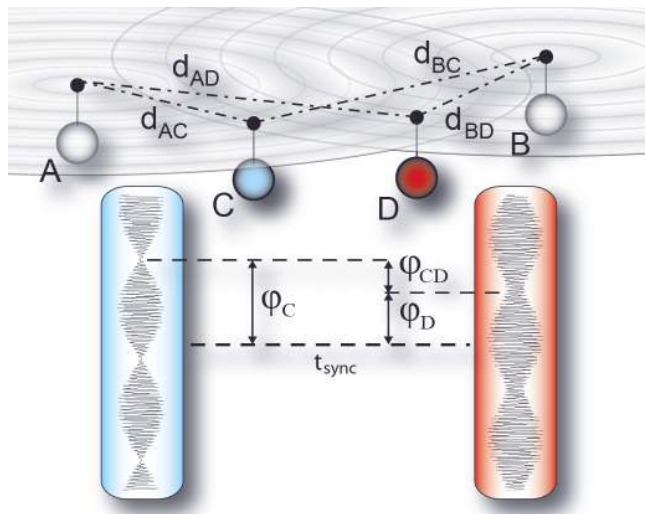


Figure 1: Nodes A and B generate an interference field by transmitting unmodulated sine waves at different frequencies. Receivers C and D measure the phase of the signal, the measured phase difference carries the location related information.

accurately. These two problems are interrelated as the generated beat frequency depends on the radio chip's ability to tune to a precise radio frequency and the frequency of the beat signal, in turn, influences the design parameters of the phase measurement algorithm. Our RIPS implementation on XSM motes benefits from the capability of Chipcon CC1000 radio chip to fine-tune its transmit frequency in 65 Hz steps. Therefore, we were able to develop a tuning algorithm which constantly optimizes the radio driver settings of the participating nodes, to maintain the interference frequency at 350 Hz. The RIPS phase measurement algorithm works with the interference signal which is sampled at 8 kHz at the RSSI pin of the radio chip. Since the frequency of the interference is low, we resorted to a simple time-domain based algorithm that performs simple filtering and calculates the average length of the signal period in a fixed window. Despite the constraint of using an 8 Mhz processor, we were able to implement an online version of this algorithm, achieving better than 0.5% frequency accuracy on average. More details about both the tuning and the phase measurement algorithms can be found in [19].

The prototype RIPS implementation on XSM nodes achieved an average localization error of 4 cm for 16 nodes covering a 120×120 m area [19].

3.2 Interferometric Tracking: inTrack

Radio-interferometric localization in sensor networks focuses mostly on stationary deployments. In particular, the solution of Eqs. (3) assumes that q -range d_{ABCD} is constant during the ranging measurement. However, with the introduction of mobility, successive phase measurements are taken at different geographic locations rendering the existing technique inapplicable. We proposed an approach to address the challenge of mobility and designed the inTrack system [12].

inTrack assumes a single mobile target node and a set of stationary infrastructure nodes located at known positions.

Note that the positions of the infrastructure nodes can be determined with sufficient accuracy using the RIPS localization algorithm. The infrastructure nodes are then responsible for both ranging and then routing the measurement results to the base station, where a data fusion algorithm calculates the location of the mobile node.

Let us assume, without loss of generality, that node A is the moving target and nodes B , C , and D are infrastructure nodes at known locations. Since the distances between B , C , and D are known, we can rearrange the q-range equation (Eq. (2)) such that the known or measurable quantities are all on the left hand side:

$$d_{ABCD} - d_{BC} + d_{BD} = d_{AD} - d_{AC}. \quad (4)$$

Let us call the left hand side of this term the **t-range** d_{ACD} . Note that it relates only three nodes, not four. Geometrically, the t-range d_{ACD} represents an arm of a hyperbola (in two dimensions) with foci C, D and with the semi-major axis of the length $d_{ACD}/2$.

Since the t-range can be computed from the measured q-range and the two known distances between the infrastructure nodes, one q-range measurement constrains the location of the target to one arm of a hyperbola. The target can be localized using the fact that hyperbolae defined by t-ranges involving the same unknown node will intersect at the target location.

3.2.1 Target as transmitter vs. target as receiver

As we discussed before, to disambiguate the measured q-range for a given pair of transmitters, we need to analyze their interference at multiple carrier frequencies. Hence, we call this multi-step process a *measurement round*. It is important to note that the number of target-related q-ranges acquired during a measurement round depends on whether the target node is a transmitter or a receiver.

In a measurement round where the target node and a designated infrastructure node are the two transmitters, and the rest of the $n - 2$ infrastructure nodes are receivers, there are $\binom{n-2}{2}$ receiver pairs yielding $\binom{n-2}{2}$ q-ranges in which the target node is involved. Assuming no two receivers have the same location, $\binom{n-2}{2}$ unique t-ranges can be calculated, though, only $n - 3$ of these t-ranges are linearly independent.

If the target node is a receiver and two designated infrastructure nodes are transmitting, there are only $n - 3$ receiver pairs involving the target node, and thus only $n - 3$ q-ranges are relevant. Moreover, it is easy to see, that after transforming q-ranges to t-ranges, all t-ranges will involve the same three nodes: the two transmitters and the target node. That is, they all define the same hyperbola. As a result, one measurement round does not yield enough information to calculate the position of the target node in the target-as-receiver case.

3.2.2 Location solver

The inTrack system uses the target-as-transmitter approach for the ranging measurements. Since a single measurement round yields $\binom{n-2}{2}$ t-ranges, the corresponding system of equations is overdetermined as the only two unknowns are the target location coordinates. This improves the robustness of the system to the measurement errors as well as helps us to reduce the need to measure interference at multiple frequencies due to the $(\text{mod } \lambda)$ -ambiguity (Eqs. (3)). The smaller number of different frequencies, in

turn, helps us to curb the measurement round duration and the latency of the tracking algorithm.

The location solver is implemented as a search algorithm that computes the position of the mobile node in a predefined search space. The q-ranges serve as inputs to the location solver; they are converted to t-ranges that constrain the location of the tracked node. The search space is subdivided into blocks (subspace of Euclidean 3D space), and for each block, it counts the number of t-ranges consistent with it. The block that has the highest count will be searched recursively. The termination condition is given as a parameter.

3.2.3 Limitations of inTrack

Apart from the intrinsic limitations of radio-interferometric ranging (RF multipath distorts phase measurements), inTrack has limited scalability with respect to its computational cost as well as poor extendibility to track multiple mobile nodes.

Computational cost. Both the q-range computation and the location solver algorithm search for extrema in a predefined search space in an exhaustive manner. The computational cost thus grows quadratically with the number of infrastructure nodes (there is potentially a q-range for each of the $O(n^2)$ pairs of infrastructure nodes) and linearly with the number of mobile nodes.

Number of targets. inTrack uses target-as-transmitter approach because it allows for acquiring more location related information per measurement. Since the target node is a transmitter, this does not scale trivially with the number of target nodes. To avoid concurrent transmissions by different targets, exclusive access to the radio channels would have to be guaranteed, increasing the latency of the method.

4. TRACKING MULTIPLE NODES

The applicability of inTrack is limited for two reasons: a) tracking multiple targets is a natural requirement in many real life scenarios, and b) it has significant hardware requirements (a PC class machine is required to compute the location), restricting the feasibility of a low-power single chip implementation. Overcoming these limitations would allow for the implementation of a low-power active badge or navigation device that works in a similar fashion to GPS, assuming a pre-deployed infrastructure in a geographically constrained area. Potential applications of such a system include locating mobile sensors, low-power alternative for (terrestrial) GPS, next generation active RFID tags, active badges for personnel tracking, etc.

As an important contribution of this paper, we propose the mTrack system, an improvement over the previously described prototype with respect to scalability and computational complexity. It supports simultaneous tracking of multiple targets and has a simple analytical location solver which, with respect to computational cost, is superior to the search algorithm employed in inTrack. In addition, by measuring Doppler shifts of the interference signal, mTrack is able to improve the localization accuracy of mobile nodes and determine the velocities of the targets simultaneously.

4.1 Approach

The main difference between inTrack and mTrack is in the roles the tracked object and the infrastructure nodes play. While in inTrack, the target is a transmitter, mTrack follows the target-as-receiver pattern. This means that only infras-

structure nodes transmit in mTrack and consequently, the number of tracked objects is not limited by channel access. Assigning the receiver role to the target, however, brings forward a series of challenges that need to be overcome.

First, it is easy to notice that the amount of measurement data collected during a measurement round is significantly less than with the target-as-transmitter approach of inTrack (see Section 3.2.1). In fact, for a given pair of transmitters and a given target, all measured t-ranges correspond to the same hyperbola. To calculate the position in a two (three) dimensional space uniquely, at least two hyperbolae (three hyperboloids) are required. Therefore, at least two (three) measurement rounds with different transmitter pairs are required to acquire the target location using mTrack.

Second, unlike inTrack, the target cannot be assumed to be stationary during the radio-interferometric measurement, since mTrack relies on multiple measurement rounds. Clearly, we need to consider the target’s speed and direction when calculating its location.

Third, the extent to which measurement errors propagate to location errors depends significantly on the location of the target with respect to the infrastructure nodes. The error propagation is the smallest when the hyperbolae intersect at an angle close to the right angle. However, when the angle of intersection is acute, small errors in the hyperbola parameters will be amplified in the calculated locations. As the location of the target is found at an intersection of a few hyperbolae, it is important that a small amplification of ranging errors is achieved for all pairs of hyperbolae. This imposes limits on the locations of the infrastructure nodes.

Finally, it is infeasible to use inTrack’s computationally expensive localization algorithm with multiple targets. We propose a simple and fast algorithm that finds the target location analytically.

4.2 Analytical Location Solver

The ranging engines of inTrack and mTrack are very similar, but the algorithms differ significantly in how they compute locations from the measured q-ranges. The basic localization algorithm of mTrack computes the location of the target node at an intersection of two given hyperbolae. Since solving this problem analytically is hard in general case, we propose to solve a special case instead, which imposes minimal constraints on the deployment topology of the infrastructure nodes. Specifically, we analyze the case when the two hyperbolae share a focus. In practice, assuming three infrastructure nodes A, B, and C, and the target X, the transmitter pair in the first measurement round would be A and B, constraining the location of X to a hyperbola with foci at A and B. In the second measurement round, nodes A and C would transmit, yielding a hyperbola with foci at A and C. Since node A always participates in the transmission, it can coordinate the measurement process for both hyperbolae and we call it a *master* node. We call the co-transmitting nodes B and C *assistant* nodes. The location of the target X is at an intersection of the two hyperbolae.

Consider Figure 2: hyperbola h_{AB} is defined by its foci A, B and the distance R_{AB} such that for any point $X \in h_{AB}$, $|AX| - |BX| = R_{AB}$. Similarly, h_{AC} is defined by the foci A, C, and the distance R_{AC} . Given the coordinates of A, B, C and the distances R_{AB}, R_{AC} , can we determine the coordinates of the intersection points of h_{AB}, h_{AC} ?




Figure 2: Hyperbolae h_{AB} and h_{AC} given by their foci A, B and A, C, respectively.

Solution:

The following equations hold for the hyperbolae:

$$\begin{aligned} \sqrt{x^2 + y^2} - \sqrt{(x - b)^2 + y^2} &= R_{AB} \\ \sqrt{x^2 + y^2} - \sqrt{(x - c_x)^2 + (y - c_y)^2} &= R_{AC} \end{aligned} \quad (5)$$

It is possible to solve this system of two equations analytically (see Appendix A) resulting in two solutions (x_1, y_1) and (x_2, y_2) . Furthermore, cutting the plane along rays AB and AC, it can be shown that there is always at most one solution on each side. That is, the solution is unique if the target is confined in the area between the two rays.

Notice that both Fig. 2 and Eqs. (5) assume that the node A is at the origin, and that node B lies on the x-axis. Given an arbitrary node deployment and a fixed reference coordinate system, this obviously will not hold. This problem, however, is easy to solve with coordinate transformation. We compute the translation matrix M_t that translates A to the origin, and the rotation matrix M_r that rotates the translated system such that the second coordinate of B be zero. We solve Eqs. (5) for the coordinates of the target in the translated coordinate system. Multiplying the resulting vector by $(M_t M_r)^{-1}$ gives the target’s coordinates in the original coordinate system.

Another straightforward relaxation of the constraints of the above solution is to allow for three dimensional infrastructure node locations (but still solving the target position in two dimensions). This approach is particularly practical when it is not possible (due to security reasons, or to avoid excessive multipath) to deploy the infrastructure nodes on the same plane where the tracked objects are expected to be moving. In a typical deployment scenario, where the tracked objects are people and the area of interest is an urban area, infrastructure nodes can be placed on rooftops and light poles.

Due to the fact that the location solver evaluates an analytical function with the measured t-ranges as arguments and the coordinates of the three infrastructure nodes as parameters, its sensitivity analysis can also be done analytically. Depending on the application scenario, we are looking for the answer for one or both of the following two questions: Given an anchor placement, how does the error in the measured t-range propagate to the resulting target coordinates? Given a user specified upper bound on the tolerable location error, what is the region in which the target can be safely tracked (assuming a given infrastructure deployment)?

The function that is used to calculate the coordinates of the target is essentially a quadratic formula depending on

the two t-ranges R_{AB} and R_{AC} . We take the partial derivatives of this function by R_{AB} and R_{AC} and evaluate the partial derivatives for a given target location thus obtaining the amplification of ranging error at the given location. Conversely, given a bound on the amplification, one can determine the region in which the amplification is below the bound. Based on the sensitivity analysis, we observe that the error amplification is less than two, if the target is located within the parallelogram defined by the points A, B, and C. Since the area of the parallelogram is maximized if the two rays AB and AC are perpendicular, we conclude that the locations of the infrastructure nodes should be chosen so that they define the right angle (or close to right angle). The location error is then less than two times the ranging error, if the target is positioned within the rectangle defined by the infrastructure nodes. See [15] for details.

4.3 Localizing Moving Nodes

Interferometric ranging as described in Section 3 assumes that the tracked object is stationary during the measurement. Otherwise, the phase difference of the interferometric signal will have a time-varying signature, rather than a constant value. If the rate of change of the phase difference is small, the ranging problem can be solved by defining the range to be an interval, rather than a single value, thus finding the location of the node in a certain region, rather than at a single point. However, the use of this method for interferometric ranging is limited due to the modulo wavelength ambiguity of the ranges. It was shown in [12] that this method can tolerate up to one half-wavelength change of the q-range during the ranging measurement which is 35 cm in our case.

For higher velocities, we propose to analyze the beat frequency of the interferometric signal to compensate for the changing q-range. It is a well known fact that moving objects measure frequencies that are Doppler shifted, depending on their speed and the direction of movement relative to the source of the measured signal. Interferometric ranging actually measures these frequency shifts with sufficient accuracy during the phase analysis of the interferometric signals. Not used in any of the distance computations for the stationary objects, these frequency shifts have been mere byproducts of the computation.

4.3.1 Doppler shifts in radio interferometry

The Doppler effects relate to the change of frequency and wavelength of a wave that is perceived by an observer moving relative to the source of the waves. For radio waves traveling at the speed of light c , the relation between the frequency f of the emitted signal and the frequency shift Δf of the observed signal is

$$\frac{\Delta f}{f} \simeq -\frac{v_O}{c} \quad (6)$$

if the observer is moving directly away from the source with speed v_O , such that $v_O \ll c$. If the observer is moving towards the source, v_O should be taken negative. Note that the Doppler effects are measurable because the frequencies we use are relatively large (400 MHz). Consequently, the observed frequency shifts due to the Doppler effect will be 1 Hz per 0.75 m/s velocity of the target.

Let X be an object at an unknown location that participates in the interferometric tracking. X proceeds by esti-




Figure 3: Doppler effect in radio interferometric ranging.

imating its distance difference from two anchors A and B (t-range d_{ABX}) which transmit sine wave signals at slightly different frequencies. Let f_a, f_b be the frequencies of these radio signals, assuming $f_a < f_b$ without loss of generality. Let \vec{v} be the velocity vector of X and let the anchors be stationary. If X was stationary, the observed frequency of the beat signal would be $f_b - f_a$. However, if X is moving, the observed beat frequency depends on the speed of X relative to A and B .

For an arbitrary node A , let v_a be defined as the dot product of the unit vector in the direction of \vec{AX} and the velocity vector \vec{v} :

$$v_a = \frac{\vec{AX}}{|\vec{AX}|} \cdot \vec{v}.$$

In other words, v_a is the length of the projection of vector \vec{v} onto the line AX , with the negative sign if the projection vector points towards the node A and the positive sign otherwise (see Fig. 3a).

The two signals transmitted by A and B interfere at X and the frequency of the resulting beat signal is:

$$f_{\text{beat}} = f_b + \Delta f_b - (f_a + \Delta f_a) = f_b - f_a - \frac{f_b}{c}v_b + \frac{f_a}{c}v_a. \quad (7)$$

In our experiments $|f_a - f_b|$ is typically smaller than 1 kHz. Consequently, we can rewrite Eq. (7) using $\lambda = \frac{c}{f_a} \simeq \frac{c}{f_b}$:

$$f_{\text{beat}} = f_b - f_a + \frac{1}{\lambda}(v_a - v_b). \quad (8)$$

Note, that the frequency difference $f_b - f_a$ and the wavelength λ are known parameters and the frequency of the beat signal f_{beat} can be measured. Therefore, we can calculate the term $v_a - v_b$ (so called **q-speed**) relating the actual speed of the moving object relative to A and B . Even though we cannot compute v_a and v_b directly, q-speed will become important in our later calculations.

4.3.2 Velocity errors compensation

Radio interferometric tracking measures t-ranges which are the distance differences $d_{XA} - d_{XB}$ of the target X from two anchor nodes A and B . It was discussed in Section 3.1 that multiple ranging measurements at different wavelengths are required to resolve the modulo wavelength ambiguity of the calculated t-range. However, the problem is that if X is moving, different ranging measurements correspond to different $d_{XA} - d_{XB}$ values. Therefore, the velocity of the object needs to be incorporated in the calculations.

Consider Figure 3b: we show here the simplest case where the target X measures phase offsets at two different carrier




Figure 4: Approximating movement errors. The target node performs two ranging measurements, at locations X and X' . We approximate $d_{AX'}$ using the relative speed v_a of nodes X and A as $d_{AX'} \simeq d_{AX} + x_a$.

frequencies to disambiguate the t-range (in practice, more than two frequencies are needed). The points X and X' that correspond to the two measurement locations are potentially far away from each other, depending on the velocity \vec{v} of X and the duration of the ranging measurement. Assuming the velocity vector \vec{v} does not change during the ranging measurement, the system of Eqs. (3) defined in Section 3 can be rewritten as follows:

$$\begin{aligned} \Delta\gamma_i &= d_{AX_i} - d_{BX_i} \pmod{\lambda_i} \\ X_i &= X + i * \vec{v} * t_M. \end{aligned} \quad (9)$$

where t_M is the time required to make one phase offset measurement and where, by abuse of notation, we identified X with the position vector of the point X .

One problem of solving Eqs. (9) is that the velocity vector \vec{v} is unknown, adding greatly to the ambiguity of the ranging solution. Although the Doppler shift analysis provides some information on the velocity \vec{v} , it does not yield \vec{v} directly.

Rather than solving Eqs. (9) for the velocity \vec{v} and the t-range $d_{AX} - d_{BX}$, we suggest to approximate d_{AX_i} and d_{BX_i} with $d_{AX} + v_a * i * t_M$ and $d_{BX} + v_b * i * t_M$ respectively. Using this approximation, we can express the measured distance differences as

$$d_{AX_i} - d_{BX_i} \simeq d_{AX} - d_{BX} + i * (v_a - v_b) * t_M.$$

The q-speed, $v_a - v_b$ can be calculated from Doppler shifts (Eq. (8)), allowing us to compute speed-compensated phase offsets $\Delta\vartheta_i$ from the measured offsets $\Delta\gamma_i$:

$$\Delta\vartheta_i = \Delta\gamma_i - (v_a - v_b) * t_M \pmod{\lambda_i}.$$

Consequently, we rewrite Eqs. (9) to

$$\Delta\vartheta_i \simeq d_{AX} - d_{BX} \pmod{\lambda_i}. \quad (10)$$

Eqs. (10) can be solved the same way as the original equations for the stationary case. However, the term $v_a - v_b$ only approximates the change of $d_{AX} - d_{BX}$ during the ranging measurement and we need to show that the error of this approximation is small.

Consider Figure 4: node X computes its location from the beat signal generated by the anchors A and B . \vec{v} is the velocity vector of the node X and \vec{v}_a is projection of \vec{v} onto AX . We analyze here the error of the approximation of the distance d_{AX} . The error of the d_{BX} approximation




Figure 5: Estimating the velocity vector v . A target node X measures two q-speeds from the Doppler shifts: $v_M - v_{A_1}$ and $v_M - v_{A_2}$. For a given estimate of v_M , the target can calculate the lines l_M, l_{A_i} and estimate its velocity vector $v = XX'$.

is analogous. Let t_M be the measurement time and $x_a = v_a * t_M$ be the distance that X travels along the direction \vec{AX} , then $d_{AX'}$ is approximated with $d_{AX} + x_a$. As seen in the figure, this approximation always underestimates the real value $d_{AX'}$ and is perfect if \vec{v} and \vec{v}_a are the same vectors. The error ε of the approximation can be expressed as $\varepsilon = d_{AX'} - d_{AX} - x_a$.

To find the maximum error of the approximation, we differentiated $\frac{\partial \varepsilon}{\partial \alpha}$ and found that the approximation error is maximal if $\cos \alpha = -\frac{d_{XX'}}{2d_{AX}}$, or, equivalently, if $d_{AX} = d_{AX'}$ (see Appendix B). Therefore, the maximum error can be calculated as

$$\varepsilon_{max} = |d_{AX} - d_{AX} - x_a| = |d_{XX'} \cos \alpha| = \frac{d_{XX'}^2}{2d_{AX}}. \quad (11)$$

In our experiments the measurement time t_M is less than 1 second. Therefore, $d_{XX'} = |\vec{v}|t_M$ is relatively small compared to d_{AX} and we conclude that the error of the approximation is relatively small.

4.3.3 Computing velocity vectors

Our localization approach calculates the position of a target node at an intersection of two or more hyperbolae defined by two or more interferometric ranges. As discussed in Section 4.2, we further require that one transmitter node, called the master node, participates in all ranging measurements. Let us recall that the co-transmitters of the master node are called assistant nodes. We denote the master node by M and the assistant nodes by $A_i, i = 1 \dots k$. We have shown that q-speeds $v_M - v_{A_i}$ can be computed from the beat frequency of the interference signal using Eq.(8). Utilizing q-speeds and the location of the target node X , we developed an algorithm that computes the velocity vector of the target node \vec{v} .

Consider Figure 5 showing a master M and two assistant nodes A_1, A_2 . The extension of our algorithm to more assistants is trivial. Let \vec{v}_M and \vec{v}_{A_i} be projections of \vec{v} , onto XM and XA_i , respectively. Line l_M is defined as the line perpendicular to \vec{v}_M which intersects the endpoint of vector \vec{v}_M placed at the target location X . Lines l_{A_i} are defined

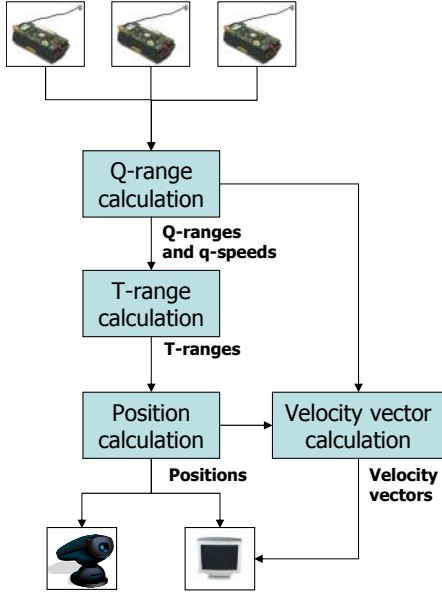


Figure 6: Dataflow structure of the mTrack system.

analogously. It is easy to see that lines l_M and l_{A_i} intersect at a single point X' such that $\vec{v} = \overrightarrow{XX'}$.

We observe, that if we know one velocity projection, for example v_M , we can compute all other projections (v_{A_i} in this case) using the measured q-speeds (see Eq. (8)). Consequently, we can find lines l_M and l_{A_i} and the intersections points of these lines P_{XY} , $X, Y \in \{M, A_i\}$. Unfortunately, v_M as well as v_{A_i} are unknown quantities as only their differences are measured.

However, it is easy to see that for incorrect v_M , the calculated v_{A_i} will all be incorrect and the intersections of the lines l_M and l_{A_i} will diverge from each other. This helps us to find the correct value of v_M . Consequently, given the maximum speed v_{\max} of the target as a parameter, our algorithm iterates through all possible values of v_M from $[-v_{\max}, v_{\max}]$ interval with certain resolution and in each iteration, calculates all intersection points P_{XY} . Ideally, for the correct value of v_M , all these intersection points correspond to a single point which is X' .

More formally, in the i -th iteration step, let v^i be the value of the speed of the target in the direction of the master node M (v_M). We define X^i as the center of gravity of all intersection points P_{XY}^i obtained using v^i as the target speed. Further, we define a quality metrics of the solution X^i as the average distance of all points P_{XY}^i from their center of gravity X^i . The final result X' of our algorithm is then such X^i which has the minimal quality metrics for all v^i from $[-v_{\max}, v_{\max}]$. Finally, the velocity vector \vec{v} is given by the vector $\overrightarrow{XX'}$.

4.4 Implementation

The mTrack system is implemented using XSM motes from Crossbow, Inc. (a variant of the Berkeley Mica2 mote) as target and infrastructure nodes, as well as a PC laptop running the location computation and the tracking GUI.

The motes are running TinyOS version 1.1.14 as the operating system. To assure that the phase measurements

are carried out at the same time on all receivers, the receivers need to be time synchronized. Instead of choosing a synchronization service that maintains a global time in the whole network (such as FTSP [20]), we opted for a multihop extension of the Estimated Time on Arrival (ETA) approach [14]: the transmitter node at the common focus of the hyperbolae, called the master node, generates a beacon event tagged with the start time of the next measurement round by the node's local clock. ETA will propagate this beacon message to all nodes within a few radio hops, converting the timestamp to the local time of the recipients. ETA is able to achieve much better utilization of system resources than virtual global time services, because it does not synchronize the clock skews of different nodes. However, the beacon event and the start of the measurement round need to be close in time (300 ms), to achieve high synchronization accuracy despite the drifts of the unsynchronized local clocks at different nodes. Notice, that it is imperative that multihop time synchronization be used, because the maximum interferometric range exceeds the communication range of the motes.

The radio-interferometric ranging engine is implemented as an alternate radio driver. Since both the TinyOS radio stack and the ranging engine use the same radio hardware, the former is disabled during the measurement rounds. After a measurement round is completed, the results (phase and frequency measurements for each frequency channel) are routed to the PC using the Directed Flood Routing Framework (DFRF) [18]. DFRF is configured with the gradient convergecast policy to provide a fast and reliable data collection service in a multihop network.

The computation of the interferometric q-ranges, their conversion to t-ranges, as well as the location computation are implemented in Java and run on a PC. The computation is implemented in a dataflow like manner, which would make it possible to distribute computational blocks to different computers, or to implement them in hardware.

The overview of the Java implementation is shown in Fig. 6. Sensors record the phase and frequency of the beat signal, while the transmitter pair iterates through a series of transmit frequencies. Q-ranges are calculated for every receiver pair and then converted to t-ranges, which are input to the analytical location solver. Once positions are known, q-speeds are converted to velocity vectors. Positions and velocities are displayed on a map and optionally used to control camera pitch/pan.

4.5 Results

Our experimental hardware platform is a Mica2 compatible XSM mote [4], programmed using the nesC programming language [5] and the TinyOS operating system [10]. Our test environment is the empty Vanderbilt Football Stadium. This allows for the direct comparison of mTrack's results with those of inTrack [12].

The mTrack system was able to get a position fix for each of the targets in approximately 4 seconds which includes 0.5 second coordination time, 1 second ranging time, 2 seconds multihop routing time, and 0.5 second localization time. During the coordination phase, infrastructure nodes were assigned the roles of transmitters and receivers and were time synchronized with each other and with the target nodes. The nodes measured the phase and frequency of the beat signal during the ranging phase and routed the mea-




Figure 7: We show anchor nodes as black dots, mobile node track by a black line, and location and velocity vector estimates by blue crosses and arrows, respectively. Map of the Vanderbilt football stadium is shown in gray. 5 anchor nodes are deployed and the mobile node starts at 45-th yard line.

sured values to the base station during the routing phase. Finally, a PC computer calculated the target locations and velocity vectors during the localization phase.

We estimated the accuracy of mTrack’s localization and velocity vector estimation with respect to the ground truth. Measuring the ground truth locations of multiple moving targets accurately, both spatially and temporally, however, was a difficult problem. We simplified this problem by pre-defining the tracks that the targets followed during the actual experiment. Each track consisted of a series of waypoints and the targets moved between two consecutive waypoints at an approximately constant speed on a straight line. During the actual experiment, we recorded the times at which each of the targets passed each of its predefined waypoints. This allowed us to compute the actual speed of the target for each segment of its predefined track. Moreover, since we also recorded the times when the ranging measurements were taken, we could determine the segment and interpolate the ground truth location of the target on that segment for any given ranging measurement. Therefore, for any given ranging measurement, we were able to reconstruct the ground truth location of the target as well as its velocity vector.

It would seem that the 4 second position fix requirement would allow us to do very rough accuracy analysis only, as the target location would change significantly in the 4 second interval. However, the actual location related information is measured only during the ranging phase of the algorithm, the beginning of which can be determined very precisely by analyzing the time-synchronization messages. Recall that




Figure 8: 4 anchor nodes are deployed. Mobile blue node follows ABCAB track, the black node is standing at C and the brown node is standing in the middle of AC.

mTrack compensates for the velocity of the targets and the target location estimate, as computed by mTrack, corresponds to the position of the target at the beginning of the ranging phase (see Section 4.3.2 and Fig. 4). Therefore, we can determine the expected location of the target accurately which allows for a meaningful error analysis of mTrack’s localization.

We conducted two experiments to illustrate tradeoffs between the accuracy and the infrastructure cost of our system. In our first experiment, we deployed five anchor nodes at known surveyed locations, covering an area of approximately 27.4×27.4 m and tracked two nodes simultaneously. We placed four anchors in the corners of this square and the fifth anchor close to the center. The actual setup can be seen in Fig. 7 which shows the anchor node locations, the Vanderbilt stadium map, the track that a mobile node follows, and the calculated locations and velocity vectors. In our second experiment, shown in Fig. 8, we decreased the number of anchors to four and tracked three nodes simultaneously in the same area. We moved the mobile node at different speeds and in different directions, collecting approximately 100 data points for each experiment. The experiments took 10 minutes to complete and we have achieved an update rate of 4.5 seconds. To be comprehensible, figures for both experiments show only a short subset of the whole dataset.

Note that in these experiments, only one node was mobile as establishing the ground truth for multiple nodes was difficult. Also, this allowed us to analyze the accuracy of our algorithm for both mobile and stationary nodes. The computational and radio bandwidth requirements of the nodes are identical, whether they are moving or not, as all of them run the same mTrack algorithm.

Three types of errors are evaluated: localization, speed, and direction error with respect to the ground truth. We calculate these errors for each of the nodes in both experiments obtaining approximately 100 data points per node per




Figure 9: Simultaneous tracking of multiple nodes: three different nodes are moving along 3 different tracks, starting at points A,B, and C.

error type. Consequently, we analyze these errors separately for mobile and stationary nodes and calculate the average as well as 95-th percentile for all error types. Table 1 shows the results achieved in both experiments, *Mobile* and *Stationary* tables showing the accuracy of mTrack only for mobile and only for stationary nodes, respectively.

MOBILE:	position (m)	Speed (m/s)	Angle (degree)
Experiment 1	0.94 (1.82)	0.18 (0.45)	11.61 (30.88)
Experiment 2	1.64 (6.41)	0.2 (0.51)	13.29 (41.84)

STATIONARY:	position (m)	Speed (m/s)	Angle (degree)
Experiment 1	0.54 (1.1)	0.19 (0.67)	N/A
Experiment 2	0.83 (1.11)	0.2 (0.65)	N/A

Table 1: Accuracy of tracking: tables show the average absolute error and the value of the 95%-th percentile in the parenthesis.

As we can see, we can achieve 0.7 m better average accuracy using more anchor nodes. Using fewer anchor nodes, our ranging data is less accurate which results in higher chance of large localization errors (95-percentile is much larger). Further, the localization error is smaller for the stationary nodes. This is due to the fact that we do not compensate the mobility related errors completely, but only use an approximation described in Section 4.3.2. The error of the speed estimation in the stationary node case is due to the occasional errors in measuring the Doppler frequency. Finally, since the ground truth for the angle measurement in the stationary case is undefined, we show it as N/A.

We also wanted to test the case when multiple mobile nodes were tracked at the same time. However, due to the difficulties in recording the ground truth in this case, we only tested our system in a series of relatively short experiments. Even though the number of collected data was not statistically significant, we have observed errors similar to




Figure 10: Simultaneous tracking of multiple nodes: a person holds two nodes in two hands, approximately 1.5 m apart and walks on the rectangular track, second person holds a single node and walks on the triangular track.

the mobile case in Table 1. Two different experiments are shown in Fig. 9 and Fig. 10.

The average localization error of the inTrack system [12] was estimated to be 0.6 m for mobile nodes. Even though our average localization error is higher, it is worth the improved scalability of our system to theoretically arbitrary number of tracked nodes as opposed to inTrack’s limit of one. The other drawback of mTrack is that we have covered a relatively small 27.4×27.4 m area. However, the mTrack system required only 4 or 5 anchor nodes which is approximately the same anchor density as the inTrack system requiring 12 anchor nodes in an 80×90 m area. Therefore, an extension of the mTrack system should also be economically viable.

5. APPLICATION EXAMPLE

To create a proof-of-concept prototype application of the radio interferometric tracking system, Vanderbilt and Oak Ridge National Laboratory teamed up to design and implement a dirty bomb detection and localization system. The operational concept is as follows. In an outdoor sports stadium or a large indoor facility, a guard walks around the stands carrying a radiation detector and a wireless node that participates in the tracking algorithm. The tracking information along with the detector readings are available in real-time to the command center. The current position is supplied also to a pan-tilt-zoom camera that automatically tracks the guard at a relatively wide angle. When the radiation level exceeds a threshold, an alarm is raised and the camera zooms in on the position of the guard. The accuracy of the system allows to zoom-in on a few meter wide area narrowing down the source to a handful of people in a crowded stadium. The guard does not even need to be aware of the alarm, in order not to raise the suspicion of the perpetrator or cause panic. The overall system architecture is shown in Fig. 11.

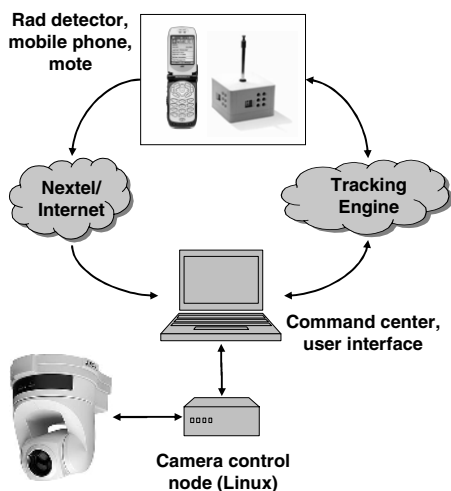


Figure 11: Dirty bomb system architecture.

The guard carries a Crossbow XSM node as well as the radiation detector connected to a mobile phone. The radiation readings are transmitted via the mobile phone network to the host computer. Hence, the host computer gets the tracking and radiation data separately. It displays the radiation level and the track on the user interface and optionally inside Google Earth utilizing a 3D model of the stadium. It also sends the position coordinates to the camera control node. When an alarm is raised, it instructs the camera to zoom in on the current position. The camera feed is accessible through a web interface.

For the proof-of-concept demonstration, radiation detection is carried out with a small, inexpensive, but highly sensitive, commercial gamma detector manufactured by RF-Trax, Inc. [24]. The crystal detector is mounted to a small circuit board that provides I/O logic and a serial interface. Detector readings are communicated to the host computer in the command center through a mobile phone and a commercial wireless TCP/IP network.

We have used an Axis 213 PTZ network camera with pan angle range -170° to 170° , tilt angle range -90° to 10° , and horizontal viewing angle of 42° to 1.7° . The camera controller accepts position commands as a point in 3D space, (x, y, z) , along with a parameter that specifies the desired radius for the radius-of-view. The former directly translates to the camera's pan and tilt settings, and the later is used to compute the camera's zoom value.

The accuracy of the camera's tracking operation is significantly affected by its initialization. For this purpose, the camera is placed at a known point $C(x_C, y_C, z_C)$ pointing at an arbitrary point-of-reference $R(x_R, y_R, z_R)$. The camera's corresponding pan θ_i and tilt α_i along with the C and R coordinates are stored in the configuration file.

To correctly slew and zoom the camera C to an object O in space with relative precision, the camera's pan angle as well as the direction of the pan has to be computed with reasonable accuracy. Given our initial pan angle θ_i of a reference point R , the pan angle θ for pointing the camera C to an object O can be computed as

$$\theta = \theta_i + \cos^{-1}\left(\frac{|CR|^2 + |CO|^2 - |RO|^2}{2|CR||CO|}\right).$$



Figure 12: The 3D model of the Vanderbilt Football Stadium in Google Earth showing the 12 infrastructure nodes (small balloons) and the guard walking on the 20-yard line (large balloon).

The tilt angle α is the slope of the vector CO : using the elevation of object O relative to the camera C ($Z_O - Z_C$), and the camera's initial tilt angle α_i , the slope is computed:

$$\alpha = \alpha_i + \tan^{-1}\left(\frac{|Z_O - Z_C|}{|CO|}\right).$$

The camera's zoom value Z is determined by the specified radius-of-view (rov) and its range of horizontal viewing angles. The zoom value for the Axis 213 camera is an integer value in the range of 1 to 9999, which corresponds to the horizontal viewing angle of 42° to 1.7° , respectively. Using these given values, we can compute the camera's angle-of-view φ and zoom value ζ as

$$\varphi = \tan^{-1}\left(\frac{2\text{rov}}{|CO|}\right)$$

$$\zeta = -169.5 + 7120\varphi.$$

The dirty bomb detection and localization prototype was successfully demonstrated in the Vanderbilt football stadium. The stadium was nearly empty, due to logistical reasons. The tracking subsystem that we used was inTrack, but mTrack fits the given scenario equally well and could be used with no changes to the integrated system. We deployed 12 infrastructure nodes, six of them on the field and the other six on one side of the stands as shown in Fig. 12. This provided a coverage of the whole grass area and half of the stands area which was approximately 80x90 meters. We estimate that to provide accurate tracking coverage in the entire 40,000 seat stadium including the field itself would require about 30 infrastructure nodes. For detailed information, including video clips of the demonstration, visit <http://www.isis.vanderbilt.edu/projects/rips/>.

The tracking accuracy was analyzed in [12]. The only area where we experienced degraded performance was in the stands above the line of the uppermost nodes. We believe that RF multipath is the culprit, as that part of the

structure is pure metal. The camera control precision, as well as the calibration accuracy, were not analyzed in detail. When pointing to positions on the stands on the other side of the stadium, translating to about a hundred-meter distance, there was a noticeable additional error when zoomed in completely. We estimated this to be under a meter or about one degree.

6. DISCUSSION

Radio interferometric ranging is able to provide high precision results utilizing low-power and low-cost hardware. Wireless sensor networks, where node localization is still an active research area, is one of the most important domains that can benefit from this technology.

The dirty bomb detection demonstration showed that inTrack, our initial radio-interferometric tracking prototype, is a feasible solution for cooperative tracking of people in stadiums and other open outdoor spaces. We expect that mTrack, the proposed extension of this system to track multiple objects, will perform equally well in real-life applications. One can envision a permanent installation of infrastructure nodes, with solar panels or with a permanent power source to avoid the need for periodic battery changes, which would enable applications such as dirty bomb detection, locating lost children in large crowds or tracking players by automatic cameras during the game.

Notice that our current system utilizes COTS hardware that was not designed for this purpose. We believe that a wireless node specifically tailored for interferometric ranging would have many benefits. It could help overcome the current speed and accuracy limitations as well as help in defeating multipath effects. Hardware support could speed up the measurements, enable a wider range of measurement frequencies and could be able to support measurements at multiple channels simultaneously. Eventually, a single-chip solution, containing a microcontroller, radio transceiver and the interferometric ranger could become an inexpensive option when mass produced.

In addition to designing a new platform, our future plans include laying down the theoretical foundations for countering RF multipath in order to make the technology available indoors and enhancing the speed and accuracy on the current platform.

7. REFERENCES

- [1] P. Bahl and V. N. Padmanabhan. Radar: An in-building RF-based user-location and tracking system. *Proc. IEEE INFOCOM*, 2:77584, Mar. 2000.
- [2] R. R. Brooks, C. Griffin, and D. S. Friedlander. Self-organized distributed sensor network entity tracking. *The International Journal of High Performance Computing Applications*, 16(3), 2002.
- [3] R. R. Brooks and P. Ramanathan. Distributed target classification and tracking in sensor networks. *in Proceedings of the IEEE*, pages 1163–1171, Aug. 2003.
- [4] P. Dutta, M. Grimmer, A. Arora, S. Bibyk, and D. Culler. Design of a wireless sensor network platform for detecting rare, random, and ephemeral events. *4th Int. Conf. on Information Processing in Sensor Networks: Special track on Platform Tools and Design Methods for Network Embedded Sensors*, Apr. 2005.
- [5] D. Gay, P. Levis, R. von Behren, M. Welsh, E. Brewer, and D. Culler. The nesC language: a holistic approach to networked embedded systems. *Proceedings of Programming Language Design and Implementation (PLDI)*, June 2003.
- [6] I. Getting. The global positioning system. *IEEE Spectrum* 30, page 3647, Dec. 1993.
- [7] L. Girod, M. Lukac, V. Trifa, and D. Estrin. The design and implementation of a self-calibrating acoustic sensing platform. *in Proc. ACM 4th Conference on Embedded Networked Sensor Systems (SenSys)*, Nov. 2006.
- [8] A. Harter, A. Hopper, P. Steggesand, A. Ward, and P. Webster. The anatomy of a context-aware application. *In Mobile Computing and Networking*, page 5968, 1999.
- [9] J. Hightower, R. Want, and G. Borriello. SpotON: An indoor 3D location sensing technology based on RF signal strength. *UW CSE 00-02-02*, Feb. 2000.
- [10] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. *in Proc. of ASPLOS 2000*, Nov. 2000.
- [11] L. Hu and D. Evans. Localization for mobile sensor networks. *International Conference on Mobile Computing and Networking (MobiCom '04)*, 2004.
- [12] B. Kusý, G. Balogh, A. Lédeczi, J. Sallai, and M. Maróti. inTrack: High precision tracking of mobile sensor nodes (accepted). *4th European Workshop on Wireless Sensor Networks (EWSN 2007)*, Jan. 2007.
- [13] B. Kusý, G. Balogh, P. Völgyesi, J. Sallai, A. Nádas, A. Lédeczi, M. Maróti, and L. Meertens. Node-density independent localization. *Information Processing in Sensor Networks (IPSN 06) SPOTS Track*, Apr. 2006.
- [14] B. Kusý, P. Dutta, P. Levis, M. Maróti, A. Lédeczi, and D. Culler. Elapsed time on arrival: a simple and versatile primitive for canonical time synchronization services. *International Journal of Ad Hoc and Ubiquitous Computing*, 2(1), 2006.
- [15] B. Kusý and J. Sallai. Analytical solution for radio-interferometric localization of mobile sensors. *Technical Report (ISIS-06-710)*, <http://www.isis.vanderbilt.edu>, Dec. 2006.
- [16] S. Lanzisera, D. T. Lin, and K. Pister. RF time of flight ranging for wireless sensor network localization. *Workshop on Intelligent Solutions in Embedded Systems (WISES'06)*, June 2006.
- [17] D. Li, K. D. Wong, Y. H. Hu, and A. M. Sayeed. Detection, classification and tracking of targets in distributed sensor networks. *IEEE Signal Processing*, 2002.
- [18] M. Maróti. Directed flood-routing framework for wireless sensor networks. In *Proceedings of the 5th ACM/IFIP/USENIX International Conference on Middleware*, pages 99–114, New York, NY, USA, 2004. Springer-Verlag New York, Inc.
- [19] M. Maróti, B. Kusý, G. Balogh, P. Völgyesi, A. Nádas, K. Molnár, S. Dóra, and A. Lédeczi. Radio interferometric geolocation. *in Proc. ACM 3rd Conference on Embedded Networked Sensor Systems (SenSys)*, Nov. 2005.
- [20] M. Maróti, B. Kusý, G. Simon, and A. Lédeczi. The flooding time synchronization protocol. *In Proc. 2nd ACM International Conference on Embedded*

Networked Sensor Systems (SenSys 2004), pages 39–49, Nov. 2004.

- [21] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. *Proceedings of ACM Sensys*, Nov. 2004.
- [22] E. Olson, J. Leonard, and S. Teller. Robust range-only beacon localization. *Proceedings of Autonomous Underwater Vehicles*, 2004.
- [23] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket location-support system. *Proc. of the Sixth Annual ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Aug. 2000.
- [24] RFTrax. <http://www.rftrax.com>.
- [25] C. Taylor, A. Rahimi, J. Bachrach, H. Shrobe, and A. Grue. Simultaneous localization, calibration and tracking in an ad hoc sensor network. *Proc. 5th Int'l Symposium on Information Processing in Sensor Networks (IPSN)*, 2006.
- [26] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, pages 99–141, 2000.
- [27] K. Whitehouse. The design of Calamari: an ad-hoc localization system for sensor networks. Master's thesis, University of California at Berkeley, 2002.
- [28] F. Zhao, J. Shin, and J. Reic. Information-driven dynamic sensor collaboration for tracking applications. *IEEE Signal Processing Magazine*, Mar. 2002.

APPENDIX

A. HYPERBOLAE

The basic building block of mTrack's localization algorithm is the analytic solver that finds the target position at the intersection of two hyperbolae. Here we show that it is possible to find a closed form formula for the intersection if the two hyperbolae share a common focus.

Consider Figure 2: hyperbola h_{AB} is defined by its foci A, B and the distance R_{AB} such that for any point $X \in h_{AB}$, $|AX| - |BX| = R_{AB}$. Similarly, h_{AC} is defined by the foci A, C , and the distance R_{AC} . Given the coordinates of A, B, C and the distances R_{AB}, R_{AC} , we want to find the intersection points of h_{AB}, h_{AC} .

The following two equations for h_{AB} and h_{AC} can be derived using the equation defining a hyperbola in 2D plane:

$$\begin{aligned}\sqrt{x^2 + y^2} - \sqrt{(x-b)^2 + y^2} &= R_{AB} \\ \sqrt{x^2 + y^2} - \sqrt{(x-c_x)^2 + (y-c_y)^2} &= R_{AC}\end{aligned}$$

Next, subtract $\sqrt{x^2 + y^2}$ and square both equations :

$$\begin{aligned}x^2 + y^2 - 2\sqrt{x^2 + y^2}R_{AB} + R_{AB}^2 &= (x-b)^2 + y^2 \\ x^2 + y^2 - 2\sqrt{x^2 + y^2}R_{AC} + R_{AC}^2 &= (x-c_x)^2 + (y-c_y)^2\end{aligned}$$

These equations can be simplified

$$\begin{aligned}\sqrt{x^2 + y^2}R_{AB} &= bx - b^2 + R_{AB}^2 \\ \sqrt{x^2 + y^2}R_{AC} &= c_x x + c_y y - c^2 + R_{AC}^2\end{aligned}$$

and equated

$$(bx - b^2 + R_{AB}^2)R_{AC} = (c_x x + c_y y - c^2 + R_{AC}^2)R_{AB}$$

Next, we can express x from this equation and substitute it into one of the original equations, essentially getting a quadratic equation with unknown y which has a closed form solution, giving us two potential solutions.

Consequently, we can find a closed formulae for both x and y , which are the intersection coordinates of the two hyperbolae h_{AB} and h_{AC} . See [15] for the derivation of complete formulae.

B. MOBILITY ERROR APPROXIMATION

In mTrack, the measurement time is approximately 1 second which means that the position of a target can change significantly during the ranging measurement. Therefore, our calculations should incorporate the velocity related location change of the targets during ranging. Due to the difficulties of determining the velocity vectors, we have only approximated it using q -speeds, defined in Section 4.3.1, which can be measured from the Doppler frequency shifts.

Consider Figure 4: if \vec{v}_A is the projection of \vec{v} onto AX , then $x_A = v_A * t_M$ where t_M is the measurement time. As noted in Section 4.3.2, we approximate distance $d_{AX'}$ by $d_{AX} + x_A$. We are interested in finding the error of this approximation which depends on the angle α (i.e. if $\alpha = 0$ then the approximation is perfect).

Let us define a function $\varepsilon(\alpha) = d_{AX'} - d_{AX} - x_a$. We want to find its extremal points. From the definition of cosine, and the law of cosines, we know that

$$\begin{aligned}x_a &= d_{XX'} \cos \alpha, \\ d_{AX'} &= \sqrt{d_{AX}^2 + d_{XX'}^2 + 2d_{AX}d_{XX'} \cos \alpha}.\end{aligned}$$

We find the extremes of $\varepsilon(\alpha)$ by differentiating it by α . We first differentiate x_a and $d_{AX'}$:

$$\begin{aligned}\frac{\partial(x_a)}{\partial(\alpha)} &= -d_{XX'} \sin \alpha, \\ \frac{\partial(d_{AX'})}{\partial(\alpha)} &= \frac{1}{2} \frac{-2d_{AX}d_{XX'} \sin \alpha}{\sqrt{d_{AX}^2 + d_{XX'}^2 + 2d_{AX}d_{XX'} \cos \alpha}} \\ &= \frac{-d_{AX}d_{XX'} \sin \alpha}{\sqrt{d_{AX}^2 + d_{XX'}^2 + 2d_{AX}d_{XX'} \cos \alpha}}\end{aligned}$$

Next we differentiate $\varepsilon(\alpha)$.

$$\begin{aligned}\frac{\partial\varepsilon(\alpha)}{\partial\alpha} &= \frac{\partial(d_{AX'})}{\partial(\alpha)} - \frac{\partial(d_{AX})}{\partial(\alpha)} - \frac{\partial(x_a)}{\partial(\alpha)} \\ &= \frac{-d_{AX}d_{XX'} \sin \alpha}{\sqrt{d_{AX}^2 + d_{XX'}^2 + 2d_{AX}d_{XX'} \cos \alpha}} + d_{XX'} \sin \alpha\end{aligned}$$

Now we let $\frac{\partial(\varepsilon(\alpha))}{\partial(\alpha)} = 0$ to find the extremes. It is easy to see that either $\sin \alpha = 0$, or:

$$d_{AX} = \sqrt{d_{AX}^2 + d_{XX'}^2 + 2d_{AX}d_{XX'} \cos \alpha}$$

from where $\cos \alpha = -d_{XX'}/(2d_{AX})$.

To summarize, $\varepsilon(\alpha)$ has extremes at (1) α is 0 or π , and at (2) $\arccos(-d_{XX'}/(2d_{AX}))$. Since $\varepsilon(\alpha)$ is 0 in the case (1), we conclude that $\varepsilon(\alpha)$ has its maximum at $\arccos(-d_{XX'}/(2d_{AX}))$. Interestingly, the approximation error happens to be maximal if $d_{AX} = d_{AX'}$.