

Node-Density Independent Localization

Branislav Kusy, Akos Ledeczki
Vanderbilt University
Nashville, TN 37212, USA
akos@isis.vanderbilt.edu

Miklos Maroti
Department of Mathematics
University of Szeged, Hungary
mmaroti@gmail.com

Lambert Meertens
Kestrel Institute
Palo Alto, CA 94304, USA
lambert@kestrel.edu

ABSTRACT

This paper presents an enhanced version of a novel radio interferometric positioning technique for node localization in wireless sensor networks that provides both high accuracy and long range simultaneously. The ranging method utilizes two transmitters emitting radio signals at almost the same frequencies. The relative location is estimated by measuring the relative phase offset of the generated interference signal at two receivers. Here, we analyze how the selection of carrier frequencies affects the precision and maximum range. Furthermore, we describe how the interplay of RF multipath and ground reflections degrades the ranging accuracy. To address these problems, we introduce a technique that continuously refines the range estimates as it converges to the localization solution. Finally, we present the results of a field experiment where our prototype achieved 4 cm average localization accuracy for a quasi-random deployment of 16 COTS motes covering the area of two football fields. The maximum range measured was 170 m, four times the observed communication range. Consequently, node deployment density is no longer constrained by the localization technique, but rather by the communication range.

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

General Terms: Algorithms, Experimentation, Theory

Keywords: Sensor Networks, Radio Interferometry, Ranging, Localization, Location-Awareness

Acknowledgments: The DARPA/IXO NEST program has supported the research described in this paper.

1. INTRODUCTION

We have recently proposed a novel approach to sensor node localization, the Radio Interferometric Positioning System (RIPS) [4]. RIPS creates a low-frequency interference signal by one pair of nodes transmitting simultaneously at close frequencies. The relative phase offset at a pair of receiver nodes is used to determine a distance measure between the transmitting and receiving nodes. Unlike tradi-

tional ranging approaches, which determine the pairwise distance between two sensor nodes, RIPS measures d_{ABCD} , a distance aggregate called the “q-range” involving four nodes: two transmitters A, B and two receivers C, D . We reported a localization accuracy of 3 cm in a 16-node setup covering an area of 324 m². We estimated the maximum range of RIPS on Mica2 motes to be 160 m, but this was not experimentally verified.

1.1 New contributions

In this paper we show that, although not straightforward, the 160 m range is indeed attainable while keeping the ranging error down to a few centimeters. This result is a significant improvement over the existing ranging solutions in wireless sensor networks (WSNs), both in terms of the accuracy and the maximum range. A 160 m range is approximately four times larger than the actual communication range when deployed on the ground. Therefore, localization no longer needs to constrain the deployment of WSNs.

Our analysis of interferometric ranging shows that it introduces significant ranging errors at large distances, which can be contributed to two major factors. The first problem is the ambiguity of the d_{ABCD} solution. Interferometric ranging computes the d_{ABCD} values from the phase offsets of the interference signal measured at two receivers C, D using the following equation:

$$d_{ABCD} \bmod \lambda = \varphi_{CD} \frac{\lambda}{2\pi},$$

where φ_{CD} is the measured relative phase offset of the receivers, and λ is the wavelength of the carrier frequency of the received signal. In general, infinitely many d_{ABCD} values solve this equation. It was shown in [4] that measuring the phase offsets using different wavelengths (λ) can eliminate the incorrect solutions. However, the particular choice of wavelengths used in [4] yields ambiguous results for q-ranges larger than 57 m (or smaller than -57 m). Here we explore how to ensure that d_{ABCD} is unique.

The second major error source is multipath radio propagation, which can distort the phase of the interference signal measured at the receivers. Multipath may become a significant error source with increasing node distances, even in the same relatively benign environment. The reason is that the ground-reflected radio signals are 180 degrees phase-shifted for small angles of incidence and travel almost the same distance as the direct line-of-sight (LOS) signal. Hence, the composite signal is attenuated and its amplitude becomes comparable to those of additional relatively weak multipath signals, causing noticeable phase deviations at the receivers.

Permission to make digital or hard copies of all or part 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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN'06, April 19–21, 2006, Nashville, Tennessee, USA.
Copyright 2006 ACM 1-59593-334-4/06/0004 ...\$5.00.

We propose a novel ranging algorithm that executes q-range estimation and localization in an interleaved and iterative manner. That is, feedback from the current localization result is used to constrain the search space of q-range estimation, then the new estimates are used in the next localization phase, thereby iteratively refining the result. We show that this technique effectively corrects ranging errors and significantly improves the localization results.

Today’s typical sensor network deployments are relatively dense, the nearest neighbor distance being at most 20 m. Consequently, the 160 m radius of the interferometric ranging can easily cover hundreds of sensor nodes introducing scalability problems for RIPS. Therefore, it is important to limit the amount of ranging data collected while ensuring that enough remains to solve the localization. First, we revisit the interesting problem of how many linearly independent interferometric ranging measurements exist for a set of n nodes and give a sharp upper bound improving the result given in [4]. Next, we present an algorithm that schedules the transmitters and receivers for the interferometric measurements. It limits the amount of acquired data and prevents uneven data sets that may otherwise result in the formation of well-localized clusters, but may not provide enough data to localize the whole network.

We present an evaluation of the improved RIPS technique based on two experiments using XSM motes [2]. First, we deployed 50 motes in an 8000 m² area with the neighbors 9 m apart. In this moderate multipath environment, we achieved a mean precision of 10 cm. In the second experiment, we deployed 16 motes in a rural area larger than two football fields. This setup demonstrated the maximum range to be 170 m, while the mean localization error was 4 cm.

We organize the remainder of the paper as follows. Section 2 revisits the theoretical background of RIPS. Section 3 describes the problems we face when increasing the maximum range. Section 4 addresses the scalability issues. We evaluate our system in Section 5 and discuss related research in Section 6. Finally, we offer our conclusions and future directions in Section 7.

2. INTERFEROMETRIC POSITIONING

Radio-based ranging techniques tend to estimate the range between two nodes from the known rate of radio signal attenuation over distance by measuring the radio signal strength (RSS) at the receiver. However, this technique is very sensitive to channel noise, reflections, interference from the environment among others. It was suggested in [4] to emit pure sine wave radio signals at two locations at slightly different frequencies. The composite radio signal has a low beat frequency and its envelope signal can be measured with low precision RF chips as the RSS Indicator (RSSI) signal (Figure 1). The phase offset of this signal depends on many factors, including the times when the transmissions were started. However, the relative phase offset between two receivers depends only on the distances between the two transmitters and two receivers and on the wavelength of the carrier signal. More formally, the following theorem was proven in [4]:

THEOREM 1. *Assume that two nodes A and B transmit pure sine waves at two frequencies $f_A > f_B$, and two other nodes C and D measure the filtered RSSI signal. If $f_A - f_B < 2$ kHz, and $d_{AC}, d_{AD}, d_{BC}, d_{BD} \leq 1$ km, then the relative*

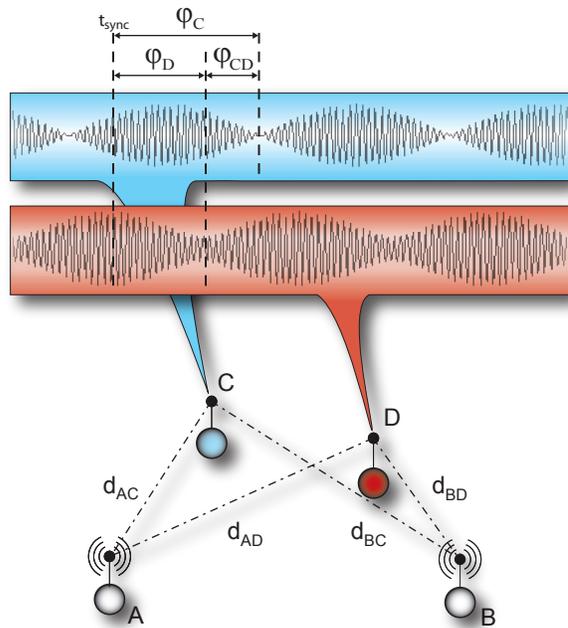


Figure 1: Radio-interferometric ranging technique.

phase offset of RSSI signals measured at C and D is

$$2\pi \frac{d_{AD} - d_{BD} + d_{BC} - d_{AC}}{c/f} \pmod{2\pi},$$

where $f = (f_A + f_B)/2$.

We call the ordered quadruple of distinct nodes A, B, C, D a **quad** and the linear combination of distances $d_{AD} - d_{BD} + d_{BC} - d_{AC}$ for quad (A, B, C, D) the **q-range** d_{ABCD} . Note that the q-range is related to the range in the traditional sense, which is the distance between two nodes, but there is a significant difference between the two measures. If d_{\max} is the maximum distance between any pair of the quad nodes, then the d_{ABCD} can be anywhere between $-2d_{\max}$ and $2d_{\max}$, depending on the positions of the four nodes.

Next, denote by φ_X the absolute phase offset of the RSSI signal measured by node X at a synchronized time instant, the relative phase offset between X and Y by $\varphi_{XY} = \varphi_X - \varphi_Y$, and the wavelength of the carrier frequency f of the radio signal by $\lambda = c/f$. Using this notation, Theorem 1 can be rewritten as

$$d_{ABCD} \pmod{\lambda} = \varphi_{CD} \frac{\lambda}{2\pi}.$$

φ_{CD} can be measured by the receivers C and D and λ is known. Note that a single φ_{CD} measurement does not yield a unique d_{ABCD} q-range because of the $(\pmod{\lambda})$ in the equation. However, we can measure φ_{CD} at different carrier frequencies, narrowing down the d_{ABCD} solution space until it contains a single q-range satisfying the maximum radio range constraint. Q-ranges can be used to determine 2D or 3D positions of the nodes, although the process is more complicated than determining positions from traditional ranges. An important difference between the interferometry and traditional ranging approaches is that we can measure at most $n(n-3)/2$ linearly independent q-ranges for a group of n

nodes, as shown in Section 4.1, as opposed to $n(n-1)/2$ linearly independent traditional pairwise ranges. Therefore, more nodes are required to determine the relative positions using interferometry. It was shown that at least 6 nodes are required to determine the 2D positions of all the nodes in the network and 8 nodes for 3D.

3. MAXIMUM RANGE OF RIPS

As discussed before, RIPS is capable of measuring the relative phase offsets of relatively weak radio signals enabling ranging well beyond the communication range. Therefore, RIPS requires multi-hop communication and time synchronization. Furthermore, the original prototype introduced in [4] may incur significant ranging errors at large distances. We analyze the sources of these errors and suggest solutions to mitigate their effects in this section.

3.1 Ambiguity of q-ranges

Denote the relative phase offset of the receivers X and Y relative to the wavelength of the carrier frequency λ as $\gamma_{XY} = \varphi_{XY} \frac{\lambda}{2\pi}$, where φ_{XY} is the phase offset of X and Y . Theorem 1 can then be restated

$$d_{ABCD} = \gamma_{CD} + n\lambda,$$

where $n \in \mathbb{Z}$, and both γ_{CD} and λ are known. Clearly, infinitely many d_{ABCD} values solve this equation (n is unknown). We can decrease the size of the d_{ABCD} solution space by measuring the phase offset γ_{CD} at m different carrier frequencies λ_i , giving γ_i , $i = 1 \dots m$. The resulting system of m equations has $m+1$ unknowns, $d_{ABCD} \in [-2d_{\max}, 2d_{\max}]$, where d_{\max} is the maximum distance between any pair of nodes, and $n_1, \dots, n_m \in \mathbb{Z}$:

$$d_{ABCD} = \gamma_i + n_i \lambda_i, \quad i = 1 \dots m \quad (1)$$

Note that this system may still have multiple solutions.

Before proceeding, we give a constraint on the n_i for later use. The phase offsets are less than 2π , so $|\gamma_i| < \lambda_i$. From $n_i = (d_{ABCD} - \gamma_i)/\lambda_i$ and $-2d_{\max} \leq d_{ABCD} \leq 2d_{\max}$, we find $|n_i| < 2d_{\max}/\lambda_i + 1$.

The problem is further complicated by measurement errors. The difference between the nominal and actual radio frequencies for a 100 ppm crystal causes an error in the wavelength of at most 0.0001λ , which can be disregarded. But the error of the absolute phase measurement on the Mica2 hardware can be as high as 0.3 rad or 0.05λ , according to our experiments, so the error of the relative phase offset γ_i can be as high as 0.1λ . We denote the maximum phase offset error with ε_{\max} and rewrite equations (1) into the following (implicit) inequalities, $i = 1 \dots m$:

$$d_{ABCD} \in [\gamma_i + n_i \lambda_i - \varepsilon_{\max}, \gamma_i + n_i \lambda_i + \varepsilon_{\max}]. \quad (2)$$

This system only has a solution if the intersection of these m intervals is non-empty. However, it is possible that the same system (with the same γ_i and λ_i values) has solutions for different assignments to the unknowns n_i , in which case the system is ambiguous. Note, that the γ_i quantities cannot be controlled. Therefore, if we want to avoid the ambiguity problem, we need to choose values for the λ_i such that ambiguity is excluded. We now derive a necessary condition for ambiguity; by contraposition, its negation is a sufficient condition for avoiding ambiguity.

So assume that also for some different vector of integers $n'_i, i = 1 \dots m$, the intersection of the intervals $[\gamma_i + n'_i \lambda_i -$

$\varepsilon_{\max}, \gamma_i + n'_i \lambda_i + \varepsilon_{\max}]$ is non-empty, and let d'_{ABCD} be a point in the resulting interval. Putting $p_i = n_i - n'_i$, we have then, $i = 1 \dots m$:

$$d'_{ABCD} - d_{ABCD} \in [p_i \lambda_i - 2\varepsilon_{\max}, p_i \lambda_i + 2\varepsilon_{\max}].$$

Since the intersection of these m intervals is non-empty, so is the intersection of any pair. This means that, for all pairs i, j in the range $1 \dots m$:

$$|p_i \lambda_i - p_j \lambda_j| \leq 4\varepsilon_{\max}. \quad (3)$$

A further constraint on the p_i values, which are integers, is found from the constraint given earlier on n_i : $|p_i| = |n_i - n'_i| < 4d_{\max}/\lambda_i + 2$. The distinctness of the vectors n_i and n'_i , finally, requires at least one p_i to be non-zero.

If, conversely, we can find λ_i such that system (3) has no solution in integers p_i —subject to the further constraints—then system (1) is guaranteed to be unambiguous for all possible outcomes for the γ_i .

We put this in context by providing concrete characteristics of our radio driver used for the Chipcon CC1000 chip: the frequency range is 400–460 MHz, the minimum separation f_{sep} between the possible frequencies is 0.527 MHz, and ε_{\max} is 0.075 m.

Within these parameters, it is actually impossible to find a set of frequencies for which (3) is unsolvable. To start, there are solutions for very small p_i . In particular, taking $p_i = 1$ for all i , insolvability requires that $\varepsilon_{\max} < \frac{1}{4} |\lambda_i - \lambda_j|$ for some pair i, j . But the range of wavelengths is 0.65–0.75 m, requiring then that $\varepsilon_{\max} < 0.025 \text{ m}$, way below the actually obtainable precision. In practice the situation is not that dire; for this to result in an actual ambiguity, all phase-measurements errors have to “conspire”, with those for the smaller wavelengths being high (positive), and those for the larger wavelengths low (negative). For a set of, say, seven wavelengths, this is rather unlikely, although not impossible. Should an ambiguity of this type occur, and should the methods explained in subsection C below not lead to the correct disambiguation, then at least the error is not extremely large.

Potentially much more pernicious are errors with large values of p_i . To express them, we rewrite system (3) into $|p_i(\lambda_i - \lambda_j) - (p_j - p_i)\lambda_j| \leq 4\varepsilon_{\max}$ and then into

$$|p_i - (p_j - p_i)\lambda_j/(\lambda_i - \lambda_j)| \leq 4\varepsilon_{\max}/(\lambda_i - \lambda_j). \quad (4)$$

Putting $d = p_j - p_i$, we have then:

$$p_i \in \left[\frac{d\lambda_j}{\lambda_i - \lambda_j} - \frac{4\varepsilon_{\max}}{\lambda_i - \lambda_j}, \frac{d\lambda_j}{\lambda_i - \lambda_j} + \frac{4\varepsilon_{\max}}{\lambda_i - \lambda_j} \right].$$

It is easy to see that all large p_i correspond to the integer multiples of $\frac{\lambda_j}{\lambda_i - \lambda_j}$, which means errors of $\frac{\lambda_i \lambda_j}{\lambda_i - \lambda_j}$ in d_{ABCD} solution space. Note, that $\lambda_i \lambda_j / (\lambda_i - \lambda_j)$ corresponds to the wavelength c/f_{sep} , where $f_{\text{sep}} = |f_i - f_j|$ is the frequency separation of f_i, f_j . We later use these two notions interchangeably. Fortunately, it is not hard to find relatively small “perfect” sets of frequencies, i.e., sets of frequencies for which such large errors are excluded.

The carrier frequencies at which the phase offsets were measured in the previous work [4] were equally distributed in the 400–460 MHz range with $f_{\text{sep}} = 5.27 \text{ MHz}$, the wavelength of which is 57 m. Therefore, to increase the range, we currently use 0.527 MHz separation with wavelength of about 569 m, allowing q-ranges up to 275 m.

3.2 Multipath effects

The results reported in [4] were obtained on a grassy area on campus near buildings and trees. As we experimented with extending the range of RIPS, the results quickly deteriorated. Once the cell size reached 10 m in the grid setup, the ranging error distribution got significantly worse. However, if we elevated the motes off the ground, the results improved markedly again. The nodes needed to be less than 10 m apart if they were on the ground, but could be more than 20 m apart if they were 4 feet high on tripods.

Figure 2 shows representative phase offset measurements on 120 channels in the 400–460 MHz frequency range with the nodes on the ground (a) and 1.3 m elevated (b). Notice that the variance of the measured phase offsets is significantly smaller in the elevated scenario, while there are severe fluctuations in case of ground deployment. When we repeated the same experiment in a rural area far from buildings and trees, the results were very close to the ideal case irrespective of the deployment height. The last observation suggests that multipath propagation is at play here, but why does elevating the motes apparently fix the problem?

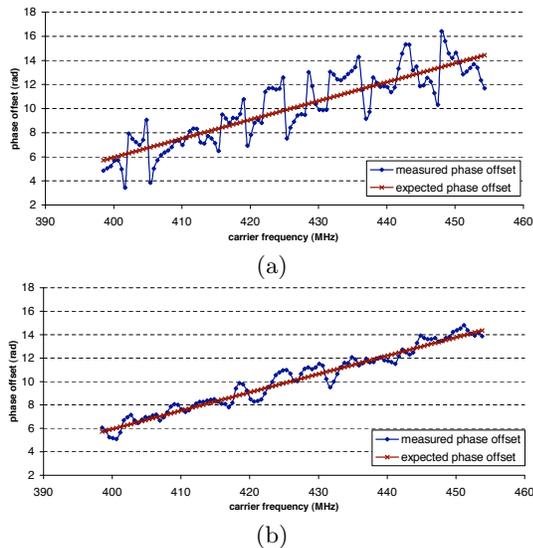


Figure 2: Phase offset measurements on 120 channels with vertical monopole antennas directly on the ground (a) and 1.3 m elevated (b).

Up till now, we have considered the radio nodes as if they were operating in free space. However, the ground around and under the antenna and other nearby objects such as trees or buildings can have significant impact on the shape and strength of the radiated pattern. These interactions can be explored in two distinctive regions surrounding the antenna. The reactive *near field* is within one quarter of the wavelength, therefore we do not consider it in this paper.

In the radiative *far field*, ground reflections—especially for vertically polarized antennas—and additional paths through reflective objects profoundly influence the received signal. When the radio wave strikes a surface, it is reflected with an angle that is equal to the angle of incidence. For surfaces with infinite conductivity, the reflected wave has the same amplitude and the same phase—or opposite, depending upon polarization—as that of the incident signal. For real surfaces, the reflected amplitude tends to be smaller and the phase relationship is more intricate. At small angles, the

phase is $-\pi$, while for larger ones, it is 0. At a certain angle, called the Pseudo-Brewster Angle (PBA) [7], the phase is $-\pi/2$. The change from $-\pi$ to 0 with increasing angle of incidence around the PBA is very steep. Thus, reflections at low angles have significant amplitude and opposite phase. As the distance difference between the line-of-sight (LOS) and the ground reflected signals is the smallest at small angles, the phase shift between them remains close to $-\pi$ and hence, the composite signal is significantly attenuated.

Figure 3 shows the simulation results of the effect of the ground reflected signal on the LOS wave over average ground surface. We used a two-node setup, where the distance between the nodes was fixed at 30 meters, while we elevated the nodes off the ground up to 5 meters—sweeping the angle of incidence between 0 and 20° . The figure shows the amplitude coefficient and the received composite signal, for which the path loss was simulated (both for the direct and reflected signals) using $\lambda^2/(4\pi d)^2$ decay. We experimentally validated the results for a smaller range of angles in a rural area where no multipath effects were present other than ground reflections and obtained data similar to the predicted values.

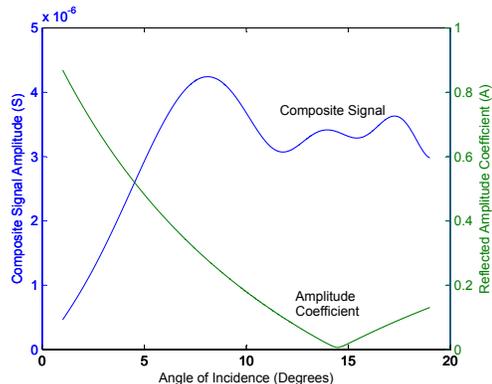


Figure 3: The reflection coefficient and the amplitude of the composite signal for vertically polarized waves. The composite signal is shaped by the interplay of the decreasing ground-reflection coefficient, and the phase offset and attenuation change due to the increasing distance the reflected signal travels.

We observed that the amplitude of the composite signal grows tenfold when we elevate the nodes from ground level to 1 m (at 30 m distance). This significant attenuation is not a problem in and of itself, as long as we can still measure the phase of the signal accurately. It definitely decreases the effective range of the method, but it does not by itself impact the accuracy noticeably. However, in a moderate multipath environment, such as the campus area we used, reflections from buildings and other surfaces distort the results. As these reflections travel longer distances, they are markedly attenuated. As long as the direct LOS signal is strong, the additional phase shift these components cause is small. However, when the ground reflection significantly attenuates the LOS signal, the phase shift caused by additional multipath components is large enough to distort the results considerably, as shown in Figure 2. We found that the measured strength of the received signals (not shown in the figure) was 8 dB stronger with elevated nodes. Since the overall topology, the node-to-node distances and the environment

were identical in the two measurements, we have experimentally verified that the differences are indeed caused by the angle—thus elevation—dependent ground reflection.

3.3 Coping with the q-range error

Intuitively, solving the ranging problem can be thought of as fitting a straight line to the measured data. As shown in Figure 2, the ideal phase offset is linear as a function of the frequency if we allow for wraparound at 2π . If we have data distorted by multipath and other errors, we can still fit a line relatively accurately, provided we have enough good data points. Therefore, a trivial enhancement is to make measurements at as many frequency channels as possible. However, this also increases the required time of the actual ranging, and a balance must be struck.

We now show how to further improve the q-range estimation and the overall localization results, even in the face of q-range ambiguity and moderate multipath effects. Let us first revisit how the baseline q-range estimation works.

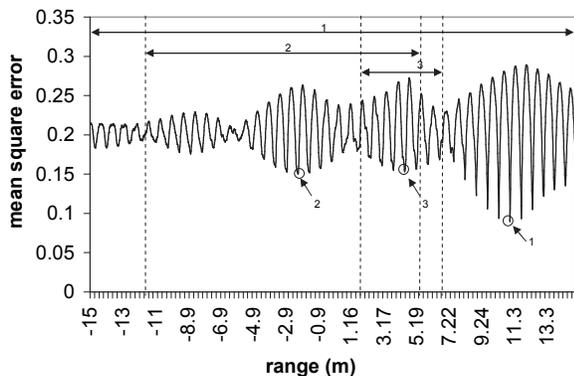


Figure 4: Phase-offset discrepancy function in a severe multipath environment. The numbers show the search intervals of consecutive error correction iterations and the corresponding minima. The spot labeled 3 is the real solution.

To get the q-range d_{ABCD} we have to solve the inequalities (2). Given a possible q-range r , for each i we can find the value of n_i that brings $\gamma_i + n_i\lambda_i$ the closest to r , namely $n_i = \text{round}\left(\frac{r-\gamma_i}{\lambda_i}\right)$. Then we define the phase-offset discrepancy function to be the average value of the squares of $\gamma_i + n_i\lambda_i - r$ values as a function of r . Ideally, the global minimum of the discrepancy function is 0, attained at the true q-range. However, measurement errors, multipath effects and the ambiguity due to the limited number of channels and the minimum channel separation, distort the results. Frequently, the global minimum is not at the true solution. Figure 4 shows an example of the phase-offset discrepancy function in a severe multipath environment indoors. The true q-range is at the local minimum labeled 3.

Given a sufficient number of q-range measurements, it is possible to estimate the node locations (up to Euclidean isometric transformations) by finding the locations that minimize the q-range discrepancy, defined as the average value of the squares of the differences between each measured q-range and the corresponding q-range “on the map”, i.e., the d_{ABCD} value computed from the node locations.

After analyzing several experiments, the following obser-

vations were made:

1. If 30–40% of the q-ranges have less than 30 cm error, the localization algorithm converges and finds approximate locations with errors smaller than a few meters.

2. Even in multipath environments the phase-offset discrepancy function has a sharp local minimum at the real q-range in most of the cases. See Figure 4.

These two observations led to the idea of the iterative phase-offset discrepancy correction algorithm:

1. Initialize the value of R , the search radius, and set all q-range estimates to 0.
2. Calculate the q-range estimates by finding the minimum of the sum of the phase-offset discrepancy functions, searching within radius R from the current q-range estimates.
3. Calculate optimal node locations.
4. If R is small enough, stop. Otherwise, decrease R and go to step 2.

In other words, the current localization solution is used to constrain the search space of the ranging algorithm, so that it can progressively eliminate large errors. Due to this feedback method, the q-range estimates get more and more accurate at each iteration.

It is easy to see that if the current value of R is always larger than the maximum q-range error in our current localization, bounding the search will not exclude the correct q-range. In our current prototype, we use fixed decreasing R values, such as $R_1 = 50$ m, $R_2 = 5$ m, $R_3 = 0.5$ m. The difference between the measured q-range and the corresponding q-range on the map is known and has a strong correlation with the localization error. This distance error could be used to drive the actual value of R and make this iterative method more adaptive. We leave this idea as a topic for future research.

4. SCALABILITY OF RANGING IN TIME

We revisit the theoretical bound on the maximum number of linearly independent q-range measurements for a set of n nodes, improving the result given in [4].

4.1 Independent q-range measurements

We assume that the network has at least three nodes, and that the nodes forming the network are numbered 0 through $n-1$. Let $N = \{0 \dots n-1\}$ denote the set of nodes. In the notation d_{AB} , we always assume that A and B are nodes in N . By convention, d_{BA} means the same as d_{AB} . Clearly, there is no need to determine quantities d_{AA} , so without loss we require in the notation d_{AB} that $A \neq B$. Then in the network there are in all $n(n-1)/2$ such quantities d_{AB} .

These distances are not independent in the sense of being mutually unconstrained. To start with, there is the triangle inequality: $d_{AC} \leq d_{AB} + d_{BC}$. Assuming that the nodes live in Euclidean 2D space, there is the further constraint that the Cayley-Menger determinant on any quad (A, B, C, D) vanishes. Here we are concerned with a more technical notion of independence: linear independence of a collection of vectors in a vector space.

Recall that, given a vector space V and a set of vectors $\{v_i\}$ of V , the subspace *spanned* by that set consists of the

collection of vectors that can be written as $\sum_i \lambda_i v_i$ for some assignment of scalar values λ_i . The set of vectors is called *linearly independent* when $\sum_i \lambda_i v_i = 0 \equiv \forall_i \lambda_i = 0$. A *basis* of V is then a linearly independent set of vectors of V that spans V .

Now take an $n(n-1)/2$ dimensional vector space over the field of the real numbers, and label the vectors of some basis with d_{AB} , for A and B distinct nodes from N — also here label d_{AB} is identified with label d_{BA} .

Define, for quad (A, B, C, D) ,

$$d_{ABCD} = d_{AD} - d_{BD} + d_{BC} - d_{AC} .$$

Thus defined, d_{ABCD} is a vector in our vector space. We call it a *measurement*, because it corresponds to a possible measurement that could be carried out by the radio-interferometric technique, the outcome being (modulo experimental error) the value of the right-hand side under some valuation of the basis vectors d_{AB} . Clearly, $d_{ABCD} + d_{BACD} = 0$ and $d_{ABCD} - d_{CDAB} = 0$, so these vectors are not all mutually independent. To rule out these pairwise dependencies, we require that in any index $ABCD$ we have:

$$A < B, \quad A < C < D, \quad B \neq C, \quad B \neq D ,$$

in which the last two inequalities, required by the distinctness of the four nodes, are given for the sake of completeness. We call an index satisfying these inequalities *normalized*. If some of the other inequalities are violated, the corresponding measurement can be found from one with a normalized index by using the pairwise dependencies given above.

Since there are three orderings of A, B, C and D compatible with the index inequalities, $A < B < C < D$, $A < C < B < D$ and $A < C < D < B$, any choice of four distinct nodes from N leads to three normalized indices, and so the set of normalized indices $ABCD$ has size $3 \binom{n}{4} = n(n-1)(n-2)(n-3)/8$.

We want to determine the dimension of the vector space spanned by the set of all possible measurements; *i.e.*, the size of a basis of that space. This is then the size of any maximally large set of linearly independent measurements.

THEOREM 2. *The dimension of the vector space spanned by the measurements d_{ABCD} on a set of n nodes, $n \geq 3$, is $n(n-3)/2$.*

PROOF. Partition the set of normalized indices into six classes:

- Class 0: $\{012D \mid 2 < D\}$ with $n-3$ elements;
- Class 1: $\{0B1D \mid 1 < B < D\}$ with $\binom{n-2}{2}$ elements;
- Class 2: $\{01CD \mid 2 < C < D\}$ with $\binom{n-3}{2}$ elements;
- Class 3: $\{0B1D \mid 1 < D < B\}$ with $\binom{n-2}{2}$ elements;
- Class 4: $\{0BCD \mid 1 < B, 1 < C < D, B \neq C, B \neq D\}$
with $3 \binom{n-2}{3}$ elements;
- Class 5: $\{ABCD \mid 0 < A < B, A < C < D, B \neq C, B \neq D\}$
with $3 \binom{n-1}{4}$ elements.

It is easily verified that all these indices are normalized and that the classes are disjoint. The sizes sum up to the cardinality of the set of all normalized indices, so this indeed constitutes a partitioning.

First we show that the measurements having indices of classes 0 and 1 together form a linearly independent set.

Next, we show that all measurements indexed by elements of classes 2-5 can be reduced to a linear combination of measurements with lower class numbers, ultimately leading to a linear combination of elements from classes 0 and 1. Combined, this gives us that classes 0 and 1 together form a basis. Since there are $n(n-3)/2$ elements in these two classes, the claim then follows.

As to the linear independence of the measurements indexed by classes 0 and 1, assume some linear combination of these measurements vanishes:

$$\sum_{2 < D} \lambda_D d_{012D} + \sum_{1 < B < D} \mu_{BD} d_{0B1D} = 0 ,$$

or, equivalently, using the definition of d_{ABCD} :

$$\begin{aligned} & \sum_{2 < D} \lambda_D (d_{0D} - d_{1D} + d_{12} - d_{02}) + \\ & \sum_{1 < B < D} \mu_{BD} (d_{0D} - d_{BD} + d_{1B} - d_{01}) = 0 . \end{aligned}$$

Recall that the vectors d_{AB} form a basis. The coefficient of each vector d_{BD} , $1 < B < D$, is $-\mu_{BD}$. So each $\mu_{BD} = 0$. Then the coefficient of each vector d_{1D} , $2 < D$, is $-\lambda_D$. So also each $\lambda_D = 0$. Therefore a linear combination of the measurements only vanishes if all coefficients are zero: they are independent.

The reductions of measurements from higher classes to classes 0 and 1 are as follows:

$$\text{Class 2: } d_{01CD} = -d_{012C} + d_{012D} ;$$

$$\text{Class 3: } d_{0B1D} = -d_{01DB} + d_{0D1B} ;$$

$$\text{Class 4: } d_{0BCD} = -d_{0B1C} + d_{0B1D} ;$$

$$\text{Class 5: } d_{ABCD} = -d_{0ACD} + d_{0BCD} .$$

In each case it is straightforward to verify that the measurements in the right-hand side are indexed by indices from a lower class. For example, the index $0D1B$ occurring in the reduction for class 3 has $D < B$ (because the left-hand side satisfies the constraints of class 3) and therefore belongs to class 1. It is equally easy to verify that each reduction represents a valid identity; for example, for class 3, expanding the definition of d_{ABCD} and using $d_{BA} = d_{AB}$, we obtain:

$$\begin{aligned} & d_{0D} - d_{BD} + d_{1B} - d_{01} = \\ & -(d_{0B} - d_{1B} + d_{1D} - d_{0D}) + (d_{0B} - d_{BD} + d_{1D} - d_{01}) . \end{aligned}$$

□

4.2 Practical solution

Measuring the q-range of all possible node quads is wasteful, as there are $O(n^4)$ quads but only $O(n^2)$ independent measurements and $O(n)$ unknowns. Furthermore, this would take an excessive amount of time in larger networks. However, trying to measure a maximal set of independent measurements is impractical when the geometry of the deployment is not known and we cannot know in advance what node quads can be measured at all. Furthermore, even successful measurements can be lost in transition. Consequently, we schedule a larger number of measurements than necessary, thereby compensating for their possible dependence, for which we do not check, while also helping to average out measurement errors.

Given a deployment, we need to select a list of transmitter pairs and the corresponding set of receivers (called a schedule) such that the collected q-range measurements are sufficient to localize. Conducting q-range measurements involving a pair of transmitters takes constant time independent of the number of receivers. Therefore, bounding the number of transmitter pairs bounds the time required to run the whole schedule. We need to consider the number of receivers that correspond to a given transmitter pair as well. Since the calculation of the q-ranges is carried out on the base station, the time required to route the phase offset measurements to the base station increases with the number of receivers, while the packet-delivery ratio decreases.

The algorithm, thus, has the following main objectives:

a) To select transmitter pairs with the most potential receivers. Since the number of linearly independent relative phase offset measurements for a given transmitter pair is $r - 1$, where r is the number of receivers, this will maximize the number of phase-offset measurements collected per transmitter pair.

b) From the set of potential receivers for each transmitter pair, to select only the best ones. This will curb the routing overhead.

c) To assure a well-connected network in terms of the node quads. This will avoid cases where some clusters of the network can be localized, but overall localization fails due to the lack of inter-cluster q-range measurements.

A simple neighborhood-discovery service provides the network connectivity as an input to the scheduler. The heuristic scheduling algorithm ranks all possible transmitter pairs by the number of neighbors the two nodes share; then picks the best ones. Next the receivers for each transmitter pair are selected based on the quality of their links from the transmitters. Coverage and connectivity is assured by best-effort heuristics: each node has to be selected as a transmitter at R_t times, and has to be a receiver at least R_r times, where R_t and R_r are empirically selected constants.

5. EVALUATION

We evaluated the improved interferometric ranging and localization in different settings. First, we had a field demonstration at the UCB Richmond Field Station. We used a 50-node approximate grid setup with a cell size of 9 m in a moderate RF multipath environment. The ground truth was obtained using differential GPS with an estimated accuracy of 1 m. 68% of the measured q-ranges had less than 1 m error, while 89% was within 2 m. Because of the inaccurate ground truth, these numbers are not revealing. However, the experiment did provide an important datapoint.

We successfully verified the performance of the scheduler because an exhaustive schedule would have taken too long, as the number of possible transmitter pairs for 50 nodes is 1225. The auto-generated schedule contained 188 transmitter pairs and 8116 d_{ABCD} measurements altogether. The actual number of q-ranges received by the base station was 4629. The 45% decrease is due to filtering and packet loss.

To get a better assessment of the overall accuracy, we hand-measured a 30-node subset of the network using measuring tape. We estimate the ground truth obtained this way to be about 5 cm accurate. The scheduler generated 107 transmitter pairs and 4517 d_{ABCD} measurements. We collected 1392 actual q-ranges from the network, of which 82% had 20 cm error or less, while 95% was better than

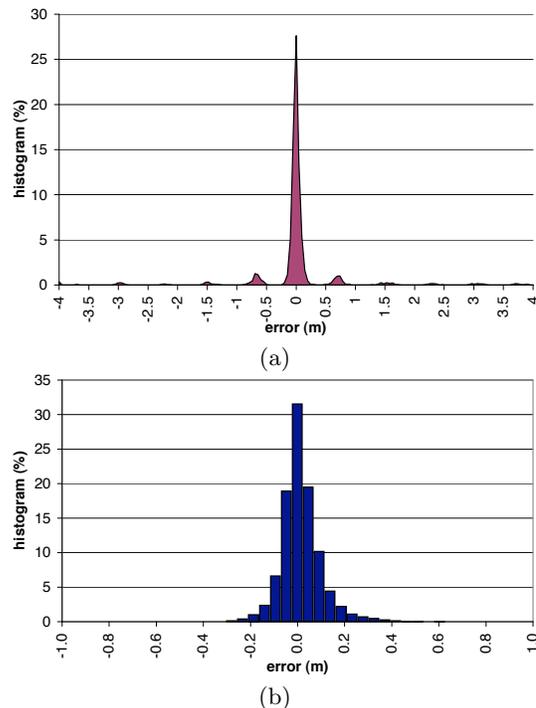


Figure 5: Error distribution of q-ranges (12000 m² experiment) obtained using (a) regular and (b) iteratively corrected ranging.

1 m accurate. The localization algorithm achieved 10 cm average accuracy, while the largest error was 20 cm.

Finally, we tested RIPS in a rural area where there were no RF multipath effects other than ground reflection, using 16 XSM motes deployed on the ground in an approximately 12000 m² area with an average closest-neighbor distance of 35 m and a maximum node distance of 170 m. We used a laser range finder to obtain the locations of the nodes with an estimated average error of 2 cm. We plot the error distribution for both non-corrected and iteratively corrected ranging in Figure 5. The original q-range error distribution had only 72% below 30 cm. Within three error-correcting iterations all q-range errors dropped below 1 m, while 98% were under 30 cm. The results of the localization from these q-ranges are shown in Figure 6.

The average localization error was 4 cm, while the largest error was 12 cm, using the three anchor nodes shown by triangles. Selecting the nodes in the four corners as anchors instead resulted in no change in the average error, but a decreased maximum error of 6 cm. The node density for this setup was 1300 nodes/km².

Using the corner nodes as anchors and keeping only four other nodes results in an 8-node setup. The subset of the measurements involving only these 8 nodes had 284 elements. This resulted in a 6 cm average and 8 cm maximum localization error, at a node density of 650 nodes/km².

6. RELATED WORK

Range-based approaches to sensor node localization provide higher accuracy than range-free methods. Ranging in WSNs is typically based on acoustic or Radio Signal Strength (RSS) measurements. Unless a powerful central

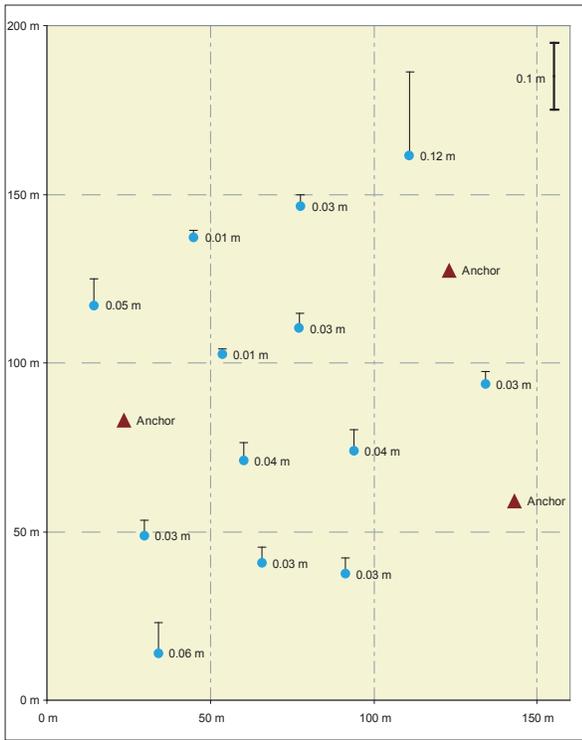


Figure 6: The setup for the 12000 m² experiment showing the anchors (large triangles), nodes (small circles) and error bars.

beacon is used as a sound source, the acoustic range is limited to 10 meters or so. The reliable range of RSS in urban deployment is even less. The largest-scale acoustic localization setup found in the literature used 45 nodes in an offset grid of approximately 9 m cell size [3]. The average accuracy achieved was 2.4 m, while the deployment density was larger than 10^4 nodes/km². Ultrasonic methods provide centimeter-scale accuracy, but due to their limited range typical setups have $> 10^5$ nodes/km² density [9, 5]. The accuracy of RSS experiments is measured in meters, with a required density of 10^5 nodes/km² [9, 1].

Recently a novel approach, called Spotlight [6], was proposed for sensor node localization. A powerful laser beacon with precisely known position and orientation is used to scan the sensor field. Nodes detect the light and record the time of detection. Correlating the detection times with the orientation of the laser, the node locations can be computed using a 3D map of the area. While the accuracy and range of Spotlight is comparable to that of RIPS, the need for an expensive central beacon, light sensors and a map of the area makes its overhead prohibitive for many deployments.

Interferometry is a widely used technique in both radio and optical astronomy to determine the precise angular position of celestial bodies as well as objects on the ground. For example, very-long-baseline interferometry, where the antennas can be thousands of miles apart, is used to track the motion of the tectonic plates on earth with millimeter accuracy [8]. However, as these sophisticated techniques require very expensive equipment, it is just the underlying idea of using radio interference for positioning that we have borrowed and applied to the sensor network domain.

7. CONCLUSIONS AND FUTURE WORK

We have successfully demonstrated in multiple field experiments that the accuracy of RIPS is as good as the best ultrasonic techniques while allowing two orders of magnitude smaller node density. At the same time, RIPS works at less than one tenth of the node density of current outdoor acoustic methods and achieves two orders of magnitude better precision.

The range that can be accurately measured with RIPS is about 4 times the communication range. That means that any deployment that forms a reasonable connected communication graph can be localized. In other words, the necessary deployment density is determined by the communication (or sensing) needs of the application and not the localization. In this sense, RIPS is indeed a node-density independent localization method.

These results correspond to minor to moderate RF multipath environments such as rural areas or open spaces within an urban area. Our current work focuses on extending the approach to cluttered urban environments.

8. ADDITIONAL AUTHORS

Gyorgy Balogh, Peter Volgyesi, Janos Sallai and Andras Nadas (Vanderbilt University).

9. 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] 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.
- [3] Y. Kwon, K. Mechitov, S. Sundresh, and G. Agha. Resilient localization for sensor networks in outdoor environments. In *25th International Conference on Distributed Computing Systems (ICDCS)*, 2005.
- [4] 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.
- [5] 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.
- [6] R. Stoleru, T. He, J. A. Stankovic, and D. Luebke. High-accuracy, low-cost localization system for wireless sensor network. in *Proc. ACM 3rd Conference on Embedded Networked Sensor Systems (SenSys)*, Nov. 2005.
- [7] R. D. Straw. *The ARRL Antenna Book, 20th Ed.* American Radio Relay League, Inc., 2003.
- [8] A. R. Thompson, J. M. Moran, and G. W. Swenson. Interferometry and synthesis in radio astronomy. *John Wiley and Sons, ISBN 0-471-25492-4*, 2001.
- [9] K. Whitehouse. The design of Calamari: an ad-hoc localization system for sensor networks. Master's thesis, University of California at Berkeley, 2002.