

Institute for Software Integrated Systems
Vanderbilt University
Nashville Tennessee 37235

TECHNICAL REPORT

TR #: ISIS-05-602

Title: Radio Interferometric Positioning

Authors: Miklos Maroti, Branislav Kusy, Gyorgy Balogh, Peter Volgyesi,
Karoly Molnar, Andras Nadas, Sebestyen Dora, Akos Ledeczki

Radio Interferometric Geolocation

Miklós Maróti^{*}
Péter Völgyesi[†]
Sebestyén Dóra[†]

Branislav Kusý[†]
András Nádas[†]
Ákos Lédeczi[†]

György Balogh[†]
Károly Molnár[†]

{miklos.maroti, branislav.kusy, akos.ledeczi}@vanderbilt.edu

ABSTRACT

We present a novel radio interference based sensor localization method for wireless sensor networks. The technique relies on a pair of nodes emitting radio waves simultaneously at slightly different frequencies. The carrier frequency of the composite signal is between the two frequencies, but has a very low frequency envelope. Neighboring nodes can measure the energy of the envelope signal as the signal strength. The relative phase offset of this signal measured at two receivers is a function of the distances between the four nodes involved and the carrier frequency. By making multiple measurements in an at least 8-node network, it is possible to reconstruct the relative location of the nodes in 3D. Our prototype implementation on the MICA2 platform yields an average localization error as small as 3 cm and a range of up to 160 meters. In addition to this high precision and long range, the other main advantage of the Radio Interferometric Positioning System (RIPS) is the fact that it does not require any sensors other than the radio used for wireless communication.

Categories and Subject Descriptors

C.2.4 [Computer-Communications Networks]: Distributed Systems; C.3 [Special-Purpose and Application-Based Systems]: Real-time and Embedded Systems; J.2 [Physical Sciences and Engineering]: Engineering

General Terms

Algorithms, Design, Experimentation, Measurement, Theory

^{*}Department of Mathematics, Vanderbilt University, 1326 Stevenson Center, Nashville, TN 37240, USA

[†]Institute for Software Integrated Systems, Vanderbilt University, 2015 Terrace Place, Nashville, TN 37203, USA

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.

SenSys'05, November 2–4, 2005, San Diego, California, USA.
Copyright 2005 ACM 1-59593-054-X/05/0011 ...\$5.00.

Keywords

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

1. INTRODUCTION

Many applications of wireless sensor networks (WSN) require the knowledge of where the individual nodes are located [1, 2, 3]. Yet robust sensor localization is still an open problem today. While there are many approaches in existence, they all have significant weaknesses that limit their applicability to real world problems. Techniques based on accurate—typically acoustic—ranging have limited range [4, 5, 6]. They need an actuator/detector pair that adds to the cost and size of the platform. Furthermore, a considerable number of applications require stealthy operation making ultrasound the only acoustic option. However, ultrasonic methods have even more limited range and directionality constraints [7, 8]. Methods utilizing the radio usually rely on the received signal strength that is relatively accurate in short ranges with extensive calibration, but imprecise beyond a few meters [8, 9, 10]. The simplest of methods deduce rough location information from the message hop count [11]. In effect, they also use the radio signal strength, but they quantize it to a single bit. Finally, many of the proposed methods work in 2D only. For a recent summary of localization methods and their performance refer to [8].

In summary, existing WSN localization methods today have either adequate accuracy or acceptable range, but not both at the same time. Furthermore, the very physical phenomenon they use—acoustics and radio signal strength—do not show any promise of achieving the significant improvement that is necessary to move beyond the current state of the art. Our novel method, on the other hand, uses radio interferometry and attains high accuracy and long range simultaneously.

Traditional radio interferometry has many applications in physics, geodesy and astronomy. The method is based on two directional antennas measuring the radio signal from a single source and performing cross correlation. The resultant interference signal can be further analyzed to create radio images of distant celestial objects, determine the relative location of two receivers very precisely, or conversely, determine the location of a radio source when the location of the two receivers are known. A radio interferometer is an expensive device requiring tunable directional antennas, very high sampling rates and high-precision time synchronization. Hence, it is not directly applicable to WSNs.

The novel idea behind the proposed Radio Interferometric Positioning System (RIPS) is to utilize two transmitters to create the interference signal directly. If the frequencies of the two emitters are almost the same then the composite signal will have a low frequency envelope that can be measured by cheap and simple hardware readily available on a WSN node. Trying to use this signal to deduce information on the positions of the two transmitters and the receiver directly would require tight synchronization of the nodes involved mandating hardware support. Instead, we use the relative phase offset of the signal at two receivers which is a function of the relative positions of the four nodes involved and the carrier frequency. By making multiple measurements in an at least 8-node network, it is possible to reconstruct the relative location of the nodes in 3D.

The key attribute of this method is that the phase offset of a low frequency signal is measured, yet it corresponds to the wavelength of the high-frequency carrier signal. Hence, we can use low precision techniques that are feasible on the highly resource constrained WSN nodes, yet we still achieve high accuracy.

The rest of the paper is organized as follows. In the next section we provide the theoretical background behind radio interferometric positioning. The subsequent section analyzes the different sources of error affecting the overall accuracy. Then we describe our prototype implementation on the MICA2 platform. It is followed by a discussion of the technique used to get a distance metric out of noisy phase offset measurements. In the subsequent section we present a centralized localization algorithm that determines the node locations from the ranging data. We conclude the paper with an analysis of the data we gathered at field experiments.




Figure 1: Radio interferometric ranging technique.

2. INTERFEROMETRIC POSITIONING

Radio interferometric positioning exploits interfering radio waves emitted from two locations at slightly different frequencies to obtain the necessary ranging information for localization. The composite radio signal has a low beat frequency and its envelope signal can be measured with low precision RF chips using the received signal strength indicator (RSSI) signal. The phase offset of this signal depends on many factors, including the time instances when the transmissions were started. However, the relative phase offset between two receivers depends only on the four distances between the two transmitters and two receivers and on the wavelength of the carrier frequency. By measuring this relative phase offset at different carrier frequencies, one can calculate a linear combination of the distances between the nodes, and ultimately infer their relative position. First we will prove these claims and then study the minimum number of measurements necessary in order to be able to resolve the phase ambiguities and localize the participating nodes.

We model the radio RSSI circuitry in the following way. The RSSI signal is the power of the incoming signal measured in dBm after it is mixed down to an intermediate frequency f_{IF} . It is then low pass filtered with cutoff frequency f_{cut} ($f_{cut} \ll f_{IF}$). Let $r(t)$ denote this filtered signal.

THEOREM 1. *Let $f_2 < f_1$ be two close carrier frequencies with $\delta = (f_1 - f_2)/2$, $\delta \ll f_2$, and $2\delta < f_{cut}$. Furthermore, assume that a node receives the radio signal*

$$s(t) = a_1 \cos(2\pi f_1 t + \varphi_1) + a_2 \cos(2\pi f_2 t + \varphi_2) + n(t),$$

where $n(t)$ is Gaussian noise. Then the filtered RSSI signal $r(t)$ is periodic with fundamental frequency $f_1 - f_2$ and absolute phase offset $\varphi_1 - \varphi_2$.

PROOF. If the noise is temporarily neglected then the mixed down intermediate frequency signal is

$$s_{IF}(t) = a_1 \cos(2\pi(f_{IF} + \delta)t + \varphi_1) + a_2 \cos(2\pi(f_{IF} - \delta)t + \varphi_2). \quad (1)$$

To obtain the signal power:

$$s_{IF}^2(t) = a_1^2 \cos^2(2\pi(f_{IF} + \delta)t + \varphi_1) + a_2^2 \cos^2(2\pi(f_{IF} - \delta)t + \varphi_2) + 2a_1 a_2 \cos(2\pi(f_{IF} + \delta)t + \varphi_1) \cos(2\pi(f_{IF} - \delta)t + \varphi_2). \quad (2)$$

Using the following trigonometric identities

$$\cos^2(x) = \frac{1}{2} + \frac{\cos(2x)}{2}$$

$$\cos(x) \cos(y) = \frac{\cos(x+y)}{2} + \frac{\cos(x-y)}{2}$$

we obtain

$$s_{IF}^2(t) = (a_1^2 + a_2^2)/2 + \frac{a_1^2}{2} \cos(4\pi(f_{IF} + \delta)t + 2\varphi_1) + \frac{a_2^2}{2} \cos(4\pi(f_{IF} - \delta)t + 2\varphi_2) + a_1 a_2 \cos(4\pi f_{IF} t + \varphi_1 + \varphi_2) + a_1 a_2 \cos(4\pi \delta t + \varphi_1 - \varphi_2) \quad (3)$$

where $(a_1^2 + a_2^2)/2$ is the DC component.

Due to the nonlinear logarithmic distortion applied to $s_{\text{IF}}^2(t)$, the resulting signal contains several new frequency components. They are the linear combinations of the frequency components $i \cdot 2\delta + j \cdot 2f_{\text{IF}}$ in equation (3), where i and j are non-negative integers.

The low pass filter eliminates all high frequency components ($j > 0$). Hence,

$$r(t) = k \cdot \log \left[(a_1^2 + a_2^2)/2 + \tilde{n}(t) + a_1 a_2 \cos(2\pi(2\delta)t + \varphi_1 - \varphi_2) \right] \quad (4)$$

where $\tilde{n}(t)$ is band-limited Gaussian white noise and k is some constant. Thus, the fundamental frequency of $r(t)$ is $2\delta = f_1 - f_2$. As the logarithm does not change the phase, its absolute phase offset is $\varphi_1 - \varphi_2$. Note that the amplitude of the harmonics are significantly smaller than that of the fundamental component. \square

We will use capital roman letters to denote nodes. The distance between nodes X and Y will be denoted by d_{XY} . The speed of light is c .

THEOREM 2. *Assume that two nodes A and B transmit pure sine waves at two close frequencies $f_A > f_B$ such that $f_A - f_B < f_{\text{cut}}$, and two other nodes C and D measure the filtered RSSI signal. Then the relative phase offset of $r_C(t)$ and $r_D(t)$ is*

$$2\pi \left(\frac{d_{AD} - d_{AC}}{c/f_A} + \frac{d_{BC} - d_{BD}}{c/f_B} \right) \pmod{2\pi}.$$

PROOF. Let X be either A or B , and Y be either C or D . Let t_X denote the time when node X starts to transmit, and a_{XY} denote the amplitude of the attenuated signal transmitted by X and received by Y . The received composite signal at node Y is

$$\begin{aligned} s_Y(t) &= a_{AY} \cos(2\pi f_A(t - t_A - d_{AY}/c)) \\ &\quad + a_{BY} \cos(2\pi f_B(t - t_B - d_{BY}/c)) + n(t) \\ &= a_{AY} \cos(2\pi f_A t - 2\pi f_A(t_A + d_{AY}/c)) \\ &\quad + a_{BY} \cos(2\pi f_B t - 2\pi f_B(t_B + d_{BY}/c)) + n(t) \end{aligned} \quad (5)$$

after sufficient amount of time, that is when t is greater than $t_A + d_{AY}/c$ and $t_B + d_{BY}/c$. Using Theorem 1, the absolute phase offset of the envelope signal $r_Y(t)$ is

$$\vartheta_Y = -2\pi f_A(t_A + d_{AY}/c) + 2\pi f_B(t_B + d_{BY}/c). \quad (6)$$

Now consider the two receivers C and D . The relative phase offset between $r_C(t)$ and $r_D(t)$ is

$$\begin{aligned} \vartheta_C - \vartheta_D &= -2\pi f_A(t_A + d_{AC}/c) + 2\pi f_B(t_B + d_{BC}/c) \\ &\quad + 2\pi f_A(t_A + d_{AD}/c) - 2\pi f_B(t_B + d_{BD}/c) \\ &= 2\pi f_A/c \cdot (d_{AD} - d_{AC}) + 2\pi f_B/c \cdot (d_{BC} - d_{BD}). \end{aligned} \quad (7)$$

From this the statement of the theorem immediately follows. \square

For wireless sensor nodes, due to their limited range and high carrier frequency relative to their cutoff frequency, the formula of the measured relative phase offset can be simplified. This gives the precursor of the definition of range following the next theorem.

THEOREM 3. *Assume that two nodes A and B transmit pure sine waves at two close 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 phase offset of $r_C(t)$ and $r_D(t)$ is*

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

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

PROOF. Using $\delta = (f_A - f_B)/2$ we can rewrite $\vartheta_C - \vartheta_D$ from the conclusion of Theorem 2 as

$$\begin{aligned} \vartheta_C - \vartheta_D &= 2\pi \frac{d_{AD} - d_{AC} + d_{BC} - d_{BD}}{c/f} \\ &\quad + 2\pi \frac{d_{AD} - d_{AC} - d_{BC} + d_{BD}}{c/\delta} \pmod{2\pi}. \end{aligned} \quad (8)$$

According to our assumption $\delta \leq 1$ kHz, so $c/\delta \geq 300$ km. Therefore, the second term can be disregarded. \square

For any four nodes A, B, C and D we define

$$d_{ABCD} = d_{AD} - d_{BD} + d_{BC} - d_{AC}, \quad (9)$$

and for any frequency f

$$\vartheta_{ABCD}(f) = 2\pi \frac{d_{ABCD}}{c/f} \pmod{2\pi}. \quad (10)$$

By the previous theorem we can effectively measure $\vartheta_{ABCD}(f)$, which equals d_{ABCD} modulo the wave length of the carrier frequency. By making multiple measurements with different carrier frequencies, one can reconstruct the value of d_{ABCD} , which is the principal ranging data of RIPS. It is trivial to verify that the following identities hold

$$d_{ABCD} = -d_{BACD} = -d_{ABDC} = d_{CDAB}. \quad (11)$$

Note that in Theorem 3 we explicitly assumed that $f_A > f_B$. If the opposite holds, then the frequency of the envelope signal becomes $|f_A - f_B|$ and its phase offset $2\pi - \vartheta_{ABCD}(f)$. This follows from the fact that $d_{ABCD} = -d_{BACD}$.

THEOREM 4. *In a network of n nodes there are at most $\frac{3}{2}(n-2)(n-3)$ independent interference measurements that can be made.*

PROOF. We fix two nodes X and Y of the network and consider the following two classes of ranges

- (M1) d_{XUYV} , where X, Y, U, V are all different nodes, and
- (M2) d_{XYUV} , where X, Y, U, V are all different nodes and $U < V$ in some fixed linear order.

Clearly, there are $(n-2)(n-3)$ ranges of the first type and one half of this number in the second type. This gives $\frac{3}{2}(n-2)(n-3)$ measurements. We claim that all other ranges can be calculated from these.

Take nodes A, B, C and D . Then it is not hard to verify that

$$\begin{aligned} &d_{XAYC} - d_{XBYC} + d_{XBYD} - d_{XAYD} \\ &= (d_{XC} - d_{AC} + d_{AY} - d_{XY}) \\ &\quad - (d_{XC} - d_{BC} + d_{BY} - d_{XY}) \\ &\quad + (d_{XD} - d_{BD} + d_{BY} - d_{XY}) \\ &\quad - (d_{XD} - d_{AD} + d_{AY} - d_{XY}) \\ &= d_{AD} - d_{BD} + d_{BC} - d_{AC} = d_{ABCD}. \end{aligned} \quad (12)$$

This proves that every range can be calculated from ranges of the form d_{XUYV} where U and V are arbitrary nodes. However, not all these are directly measurable, because for example X cannot be the same node as U or V . The two degenerate cases are

$$\begin{aligned} d_{XXYU} &= d_{XU} - d_{XU} + d_{XY} - d_{XY} = 0, \quad \text{and} \\ d_{XUYU} &= d_{XY} - d_{UY} + d_{UY} - d_{XY} = 0. \end{aligned} \quad (13)$$

Therefore, d_{ABCD} can be calculated from ranges from class (M1) whenever $X \neq C, D$ and $Y \neq A, B$.

Using the equation $d_{ABCD} = d_{CDAB}$, we can rewrite equation (12) as

$$d_{XCYA} - d_{XCYB} + d_{XDYB} - d_{XDYA} = d_{ABCD}, \quad (14)$$

which allows to calculate d_{ABCD} whenever $X \neq A, B$ and $Y \neq C, D$. If neither (12) nor (14) can be used to calculate d_{ABCD} then it must be the case that $\{X, Y\} = \{A, B\}$ or $\{X, Y\} = \{C, D\}$. In these cases, however, d_{ABCD} can be directly obtained from a measurement in class (M2) using equations (11). \square

Note that the previous theorem gives only an upper bound on the number of linearly independent set of measurable ranges in an n -node network. Since the submission of our paper, Lambert Meertens significantly improved this bound to $n(n-3)/2$, and proved that to be sharp [12].

Notice that any solution to the resulting system of equations is invariant under translations, rotations and reflection. Therefore, the number of unknowns is $2n-3$ in 2D and $3n-6$ in 3D. In a 6-node network there are 9 linearly independent measurements, using the sharp bound, just enough for the 9 unknowns. However, we need at least 8 nodes in 3D to get more measurements (20) than the number of unknowns (18).

3. SOURCES OF ERROR

The goal of this section is to list and briefly discuss the sources of error of RIPS. In most cases further research is necessary to properly study these errors.

We will consider two nodes transmitting unmodulated radio waves at two close frequencies and two receivers measuring the absolute phase offset of the received signal $r(t)$. We assume that the absolute phase offset is measured at a fixed time instant that is established between the receivers using some form of time synchronization. The relative phase offset is then calculated by subtracting the absolute phase offsets of the two receivers.

The sources of error are the following:

Carrier frequency inaccuracy—the difference between the nominal and the actual carrier frequency of the transmitted signal. According to Theorem 3, the phase measurement error introduced by an average 1 kHz carrier frequency inaccuracy for ranges d_{ABCD} less than 1 km is

$$\left| 2\pi \frac{d_{ABCD}}{c/f} - 2\pi \frac{d_{ABCD}}{c/(f+1000)} \right| \leq \frac{2\pi \cdot 1 \text{ km}}{c/1 \text{ kHz}} = 0.33\% \cdot 2\pi,$$

independently of the value of f .

Carrier frequency drift and phase noise—the phase noise and drift of the actual carrier frequency of the transmitted signal during the measurement. Theorem 1 relies on the fact that the frequencies of the emitted signals are stable. Any phase noise or carrier frequency drift will be directly

observable in the measured phase offset of the envelope signal. This source of error can be minimized by shortening the length of a single phase measurement, and by minimizing the chance of mechanical, electrical and other kinds of shocks the RF chip is subjected to. In our implementation one measurement lasts for 29 ms so frequency drift is likely negligible. Phase noise, on the other hand, may have a significant effect on phase measurement accuracy.

Multipath effects—the RF signal takes different paths when propagating from a transmitter to a receiver, causing amplitude and phase fluctuations in the received signal. In Theorem 3 we assumed that the radio signal travels from A to C for exactly d_{AC}/c seconds, but this assumption does not hold in the presence of multipath fading. We expect that in many cases higher level algorithms can filter out inconsistent phase offset measurements corrupted by multipath effects.

Antenna orientation—the change of the time of flight from the transmitter to the receiver introduced by the different orientation or shape of the antennas used. Regardless of the orientation of antennas, the received radio signal component in Equation (5) originating from A and received by C is always

$$a_{AC} \cos(2\pi f_A(t - t_A - d_{AC}/c)).$$

The frequency f_A and transmission time t_A are constants, and by Theorem 1 the amplitude a_{AC} of the signal does not influence the phase offset of the signal strength. It is theoretically possible that the time of flight component d_{AC}/c is influenced by the orientation, but this was not empirically verified.

RSSI measurement delay jitter—the jitter of the delay between the antenna receiving the radio signal and the RF chip delivering the RSSI signal to the signal processing unit. This jitter introduce a relative phase offset error at the two receivers. According to our experiments, the jitter is not noticeable.

RSSI Signal-to-Noise Ratio—the signal strength relative to noise of the $r_C(t)$ signal. The SNR mainly depends on the physical distance between transmitters and receiver, as the amplitudes of the transmitted signals are exponentially decreasing in space. The SNR value also depends strongly on the hardware implementation of the RSSI detector circuit at the receiver.

Signal processing error—the error introduced by the signal processing algorithm that calculates the phase offset of $r_C(t)$. This is a noisy, logarithmically distorted, low frequency sinusoid, that can be approximated by a sine wave with additive noise. The frequency and phase estimation of sine waves is a well-studied problem (see e.g. [13, 14]). The theoretical Cramér-Rao bound can be approximated by various signal processing algorithms for a given SNR.

Time synchronization error—the time synchronization error of the time instance when the receivers measure their absolute phase offset of the received signal strength. On representative hardware it is possible to establish a time synchronization point with better than $2 \mu\text{s}$ precision utilizing a single radio message (see [15]). Assuming a 2 kHz interference frequency, this translates to $\%0.4 \cdot 2\pi$ phase offset error.

4. IMPLEMENTATION

We have used the MICA2 mote platform with the TinyOS operating system [16, 17] for our prototype implementation. As stated in Section 2, the primary objective is to have two nodes transmit sinusoid waves at close frequencies to produce the interference signal and at least two receivers to calculate the phase offset difference of the observed signals. It is necessary to measure phase offset differences at multiple frequencies to be able to calculate the range. Specifically, our prototype RIPS implementation consists of the following steps:

- (1) selecting a pair of transmitters from a group of motes participating in the localization and scheduling their transmission times,
- (2) fine-grain calibration of the radios of senders to transmit at close frequencies,
- (3) transmission of a pure sine wave by the two senders at multiple frequencies,
- (4) analysis of the RSSI samples of the interference signal at each of the receivers to estimate the frequency and phase offset of the signal,
- (5) calculation of the actual d_{ABCD} range from the measured relative phase offsets for each pair of receivers, and
- (6) the localization algorithm.

Currently, the selection and scheduling of the transmitting pairs, part of the frequency calibration process, the calculation of the d_{ABCD} range from multiple phase offset readings and the localization itself are running on the base station. The software components running on the motes include a custom driver for the Chipcon CC1000 radio that allows pure sine wave transmission at a particular frequency, a radio engine component that coordinates and synchronizes the participating nodes and handles the transmission and reception of the interference signal, and a signal processing component that estimates the frequency and phase offset of the sampled RSSI signal.

4.1 CC1000 characteristics

Our sensor nodes are equipped with the Chipcon CC1000 radio chip configured to transmit in the 433 MHz frequency band. The following features provided by the radio chip were essential to the implementation of our interferometric algorithm:

- (1) capability to transmit an unmodulated sine wave in a reasonably wide frequency band (between 400 MHz and 460 MHz), as well as the ability to tune the frequency of the transmitter in fine-grain steps (65 Hz),
- (2) short-term stability of the frequency of the unmodulated sine wave (for less than 29 ms time period),
- (3) relatively precise capture of RSSI with only a small delay jitter, and
- (4) capability to transmit at different power levels.

The relatively wide frequency band is necessary for calculating the actual range from the phase offset differences (see Section 5). The fine-grained frequency tuning is needed to achieve the separation of the two transmitters as required by Theorem 3 in Section 2. We require the short-term stability of the carrier frequency because the frequency and phase offset analysis at the receivers uses averaging over multiple periods to increase the SNR. The timing precision of the measured RSSI signal is also critical because the RSSI measurement delay affects the relative phase offset error as discussed in Section 3. Finally, transmission at different power levels is required since the distance of the two transmitters from a receiver can vary up to the point where the closer transmitter completely overwhelms the signal from the more distant one.

Even though the radio chip is highly configurable with many favorable properties, it has certain limitations that we had to overcome:

- (1) the time required to calibrate the radio chip has a large jitter of up to 5 ms,
- (2) the frequency synthesizer is not accurate enough, we have seen carrier frequency deviations that depended on the temperature and voltage levels,

We cope with the radio timing fluctuations by using a time synchronization scheme, which is external to the radio chip (Section 4.2). It is also important to overcome the carrier frequency inaccuracy because the interference frequency needs to fall within a small range. We implemented an algorithm (Section 4.3) that determines the radio parameters for the transmitters corresponding to a frequency separation of close to 0 Hz. We run this algorithm frequently (i.e., once for each transmitter pair) to overcome the dependency of the generated frequency on the current temperature and voltage.

4.2 Time synchronization

In order to measure the relative phase offset of the RSSI signals between different receivers, the nodes need to synchronize and measure the absolute phase offsets relative to a common time instant. After collecting the absolute phase offsets from the receivers, the relative phase offsets can be calculated by subtracting them from each other. In this section we discuss the necessary synchronization strategy. We do not use network-wide time synchronization. Instead, we synchronize the nodes participating in the current ranging round only and only for the duration of a single measurement.

Before a node can transmit or receive a radio signal at a particular frequency, it needs to acquire the radio chip from the standard MAC layer and calibrate it to that frequency. Once the chip is acquired, no further inter-node communication is possible, so the participating nodes have to follow a predefined schedule starting from a precise time instant to stay synchronized. This scheduling is challenging because the time required to perform most operations with the radio, such as acquire and calibrate, have a significant variance. The crucial operation is the sampling of the RSSI signal that needs to be aligned with a couple of microseconds precision across the receivers.

The timing uncertainties in the radio chip are mitigated by imposing an external synchronization protocol depicted




Figure 2: The synchronization schedule that aligns the start of the transmission and reception at multiple nodes.

in Figure 2. In this protocol, the master node initiates the measurement by broadcasting a radio message, which identifies the other sender node, the type of measurement (tuning or ranging) and the transmit powers. The message also specifies a time instant in the future—in the local time of the master—when the measurement must be started. Finally, the radio stack attaches a precise timestamp to the packet, when it is sent from the mote. Each receiver converts the synchronization point to its local time by using the arrival timestamp of this message. The converted value is used to set up a local timer and is re-broadcasted. This simple protocol enabled us to extend interferometric measurements beyond the data communication range. It was shown in [15] that it is possible to set up synchronization points this way with microsecond accuracy on the MICA2 platform. We start calibration and signal transmission/reception at fixed times after the synchronization point. The combined error of synchronization and clock skew over a time period of less than 1 second is still just a few microseconds.

The external synchronization ensures that the receivers and transmitters send the receive and send commands to the radio chip at the same time. This does not necessarily mean that the receivers start receiving data from the ADC at the same time. To compensate for this, each receiver records the times when the RSSI samples arrived and adjusts the measured absolute phase accordingly.

4.3 Tuning

The CC1000 radio chip needs to perform internal calibration of the frequency synthesizer PLL to adjust it to the frequency and compensate for supply voltage and temperature variations. The self-calibration is a time consuming process that takes up to 34 ms according to the Chipcon datasheet [18]. Moreover, since the PLL can generate only a limited scale of frequencies, it is advised to recalibrate it for frequencies more than 1 MHz apart. We found it useful to index the available frequency band (400 MHz to 460 MHz) by frequency *channels*, so that transmitting on a different channel mandates recalibration. We define channel 0 to be at 430.1 MHz and the channel separation to be approximately 0.526 MHz.

A key benefit of CC1000 chip is that the PLL can generate different frequencies at a very fine frequency resolution (65 Hz) once calibrated to a particular channel. We index these fine-grained frequencies with frequency *tuning* and require no recalibration when changing the tuning parameter. The nominal frequency f can be then obtained from the

channel and tuning parameters the following way:

$$f = 430.1 + 0.526 \cdot \text{channel} + 65 \cdot 10^{-6} \cdot \text{tuning}.$$

One of the limitations of the chip is that the actual frequency for a specific channel can differ from the nominal value by up to 2 kHz. Since the frequency range where we can measure the interference accurately is limited by the time constraint for a measurement (29 ms) as well as the limited sampling rate of the RSSI signal on MICA2 platform (9 kHz), we need to ensure that the difference between the transmission frequencies is in the range of 200 Hz to 800 Hz.




Figure 3: A receiving node observes the frequency of interference of two transmitters. The first transmitter changes the carrier frequency in 325 Hz steps.

We implemented a frequency tuning algorithm that determines the radio settings for the transmitters corresponding to the same frequency. Let the transmitters be set to the same channel and f_1, f_2 be the actual signal frequencies. We found the maximum carrier frequency error to be less than 2 kHz, so we can assume that $|f_1 - f_2| < 2 \cdot 2 \text{ kHz} = 4 \text{ kHz}$. The first transmitter emits a sine wave at frequencies $f_1(i) = f_1 + i \cdot 325 \text{ Hz}$, $i = -15, -14, \dots, 15$ using the tuning capability of the radio chip, while the second transmitter keeps transmitting at frequency f_2 . A receiver node analyzes the frequency of the interference signal which is $|f_1(i) - f_2|$ (see Figure 3). Using the known step size (325 Hz), the receiver can filter out the noise and faulty frequency measurements and determine the value of i for which the interference frequency is close to 0. The receiver propagates this information back to the first transmitter who consequently determines the settings for the radio chip such that the interference frequency is in the required range.

It should be noted that the frequency error for two nodes does not stay constant for different channels. Observing it at two channels 50 MHz apart, we have seen up to 300 Hz change. However, as the frequency error is mainly caused by the imprecision of the crystal that drives the radio chip, it is approximately linear in frequency. Therefore, we can measure the correct tuning parameter for two different channels and interpolate these values to obtain the radio parameters for the other channels.

4.4 Frequency and phase estimation

Due to the limited communication bandwidth the sampled RSSI values need to be processed on the motes. The signal processing algorithm estimates the frequency and the phase of the RSSI signal and transmits these and a quality indicator of the measurement to the base station. The algorithm is divided to online and post-processing parts. The

online part is executed upon each A/D converter interrupt for 256 consecutive samples. Afterwards, more extensive post-processing is performed on the data computed in the online phase.

Hardware limitations on the mote make computationally expensive signal processing techniques prohibitive. The ADC sampling rate (9 kHz) and the clock frequency of the 8 bit microcontroller (7.4 MHz), allows roughly 820 CPU cycles per sample for online processing. Post processing is limited by a somewhat less strict deadline: several measurements are made between time synchronization points, leaving around 10,000 cycles per measurement for post-processing. The lack of floating point hardware support and memory space limitations further restrict the domain of feasible algorithms. The use of standard, but computationally expensive solutions, such as Fourier analysis or autocorrelation, is not feasible.

Figure 4 shows a representative RSSI signal recorded by a mote. Peak detection is performed on-line in the ADC interrupt routine eliminating the need for large sample buffers and shortening the post-processing time. First, the raw samples are filtered by a moving average component in order to enhance the SNR. Next the minimum and maximum signal values—essential parameters to our adaptive peak detection algorithm—are acquired from the leading 24 samples. This first part of the data series is long enough to contain at least one full period. The acquired amplitude value serves as a quality indicator of the measurement. Later on, samples above (below) a threshold currently set at 20% of the amplitude from the maximum (minimum) value are identified as high (low) amplitudes in the filtered signal. Peaks are defined as center points of two consecutive high level threshold crossings (non-high \rightarrow high, followed by a high \rightarrow non-high step). Peaks are discarded in this phase if the signal has not crossed the low threshold since the last peak, minimizing the risk of false positive detections.




Figure 4: Peak detection and filtering.

The post-processing phase works exclusively on peak indexes, identified and stored by the online algorithm. After it determines the shortest period between subsequent peaks, it accumulates the sum of the periods that are not longer than 130% of the shortest one (to compensate for the very rare false positive detections in the first phase). Peaks on both ends of an outlier period are marked as false peaks and discarded. This simple and draconian rule might throw out

good peaks, a small price to pay for rejecting false peaks, which could significantly impair the phase calculation. Frequency is calculated as the reciprocal of the average period length.

The phase of the RSSI signal is estimated by the average phase of the filtered peaks. Since small errors in the frequency estimation can result in a significant error in the phase calculation, we compute the phases relative to the center of the sample buffer, thereby reducing the accumulated phase error due to an inaccurate frequency estimate [13]. The algorithm also employs a basic phase unwrapping method to average values near 0 and 2π properly.

Since floating point calculations are prohibitive on the mote, hand optimized fixed point arithmetic is used throughout the frequency and phase computations. The estimated frequency, phase and amplitude tuple is finally sent back to the base station.

4.5 Scheduling

There are two levels of scheduling involved in the interference measurement process:

- A) High level scheduling is responsible for selecting the pair of transmitters and should minimize the number of interference measurements while producing enough independent measurements to localize nodes uniquely in 3-dimensions. As given at the end of Section 2, for a group of n nodes that form a single hop network, we have at most $n(n-3)/2$ choices for the independent interference measurements. The number of unknowns is $3n-6$ in 3-dimensional localization, so for groups of nodes larger than 8 we get an over-determined system of equations. In our current implementation, the base station selects all possible pairs of transmitters while all other nodes within their range act as receivers.
- B) Low level scheduling coordinates the activities of the two transmitters and multiple receivers. The frequency tuning algorithm and the phase offset estimation process described earlier in this section both involve multiple steps that require proper frequency calibration and timing. Currently, 13 different frequency channels 5 MHz apart are used between 400 and 460 MHz. Furthermore, the scheduler executes the phase offset measurement with the same pair of transmitters, but different radio power settings to compensate for the effect of one transmitter being much closer to a receiver than the other. Currently, three combinations are used: full power/full power or the two combinations of full power/half power.

5. RANGE CALCULATION

From a set of phase measurements for the frequencies f_1, \dots, f_k the following Diophantine equations can be formulated

$$d_{ABCD} = \lambda_i n_i + \gamma_i = \lambda_j n_j + \gamma_j, \quad (15)$$

where $\lambda_i = c/f_i$ is the wave length, $\gamma_i = \lambda_i \frac{\vartheta_i}{2\pi}$ is the phase offset relative to the wave length, ϑ_i is the measured phase offset, and n_i is an integer. We need a set of λ_i 's so that their least common multiple is larger than the possible domain of d_{ABCD} . In case of the 433 MHz band, having 5 MHz separation means that three different measurements are enough

assuming perfect phase estimation. The concept is illustrated in Figure 5.




Figure 5: Calculating the range from phase offsets.

Unfortunately, the phase measurements are noisy, so the Diophantine equations become invalid. They can be reformulated into inequalities

$$|(\lambda_i n_i + \gamma_i) - (\lambda_j n_j + \gamma_j)| < \varepsilon, \quad (16)$$

where ε is a fraction of the wavelength determined by the phase measurement accuracy. This inequality set can be solved and the d_{ABCD} solution defined as the mean of the individual $d_i = \lambda_i n_i + \gamma_i$ values. However, there may be multiple solutions resulting in more than one d_{ABCD} values differing by small integer multiples of the wavelength approximately.

An error function can be defined as

$$error = \sqrt{\sum (d_{ABCD} - d_i)^2}. \quad (17)$$

The solution with the minimum error value becomes the final d_{ABCD} estimate. The more number of different frequencies are used the better the estimate is. Instead of the necessary three, we currently use around 10 different frequencies.

6. LOCALIZATION

Since the RIPS ranging method does not provide range estimates between a pair of nodes directly, but a combination of distances among four nodes, none of the existing localization methods is directly applicable. Solving the large number of nonlinear equations would be somewhat cumbersome and not scale well. Instead, as a first cut at the localization problem, we decided to use an optimization method based on genetic algorithms (GA). GA follows the idea of Darwinian evolution and widely used as a general function optimization method. Our current approach is not meant as a comprehensive solution yet. In the limited time we had, our goal was to evaluate the ranging method in the context of overall localization accuracy and provide a baseline localization method.

Given a set of nodes with unknown locations and a set \mathcal{M} of d_{ABCD} ranges, our goal is to find the relative positions of the nodes. An error function is defined over the node locations and a GA is used to find the node locations with the smallest error. Here is the brief description of the applied genetic algorithm:

- (1) Generate an initial population of *populationSize* random solutions.
- (2) Select *subpopulationSize* solutions randomly from the population.
- (3) Evaluate each solution in the selected subset using the error function.

- (4) Sort the subset according to error.
- (5) Remove the worst 20% of the individuals in the subset, then generate new individuals by selecting random parents from the best 20% and applying genetic operators on the parents.
- (6) Go to step (2).

A solution here is a placement of nodes. The error of a solution s is defined as

$$error(s) = \frac{1}{n} \sqrt{\sum_{ABCD \in \mathcal{M}} (d_{ABCD} - d_{ABCD}(s))^2}, \quad (18)$$

where d_{ABCD} is the measured range and $d_{ABCD}(s)$ is the calculated range in the solution s .

A node placement is represented directly by a vector of (x, y, z) coordinates of the nodes. The following genetic operators are used to create new solutions:

- (1) Crossing over: each node position is inherited from one of the two parents with even chance.
- (2) Mutations (all cases have equal chance)
 - (a) Move one node by a Gaussian random number with a small ε variance.
 - (b) Move one node to a random position.
 - (c) Move all nodes by the same Gaussian random number with a small ε variance.

The value of ε is set to the current value of the error function. It makes it possible for the nodes to do bigger ‘‘jumps’’ if the error is larger. When the error gets small the nodes can fine tune their positions this way.

This algorithm uses all the given ranges and tries to minimize the difference between the input range and the range in the solution. However, the input data has range estimates with relatively large errors corresponding to integer multiples of the wavelength distorting the solution. We extended the genetic representation of the solution to include the set of used measurements and let the GA optimize this set as well. This way the GA searches for the node positions and a set of good measurements simultaneously, making it possible to eliminate all or most of the bad measurements. In our experiments this enhanced GA was able to reach the same accuracy as the one running on a data set where all ranges with large errors were removed manually.

7. EVALUATION

7.1 Effective range

Determining the effective range of the radio interferometric ranging technique is not as straightforward as it is with methods relying on direct pairwise ranging. There are four nodes involved here and not only are there constraints on their pairwise distances, but also restrictions on the geometry of their arrangement. The maximum distance between a transmitter and a receiver is clearly related to the radio range. We have observed nice interference signals even when the receiver was far beyond the communication range of the transmitters at 160 meters using somewhat elevated notes. In the same setup the communication range was only

80 meters. That means that the interferometric technique can have twice the range as the digital communication implemented on the MICA2 mote using the same radio. In the remaining discussion we'll call this distance the *interferometric radio range*.

So far we have only talked about the maximum distance between a transmitter and a receiver. There is no direct constraint on the distance between the two transmitters or the distance between the two receivers. However, by implication, they need to be within twice the interferometric radio range. Another important consideration is the ratio of the distances from a single receiver to the two transmitters. The received signal strength of one of the signals cannot be much larger than the other to generate a good quality interference signal. Tuning the transmitter output power can compensate for this, however. Note that this does not constrain the range of the given geometric arrangement of two transmitters and a receiver, since the given transmitter has to be close to the receiver if its signal can overwhelm the other signal. On the other hand, the same amplitude tuning needs to work for two receivers *at the same time*. Therefore, the second receiver needs to be in an area where the interference signal quality using the given amplitude tuning is still good enough. Furthermore, obstructions, multipath effects and other environmental conditions will adversely affect the effective range.

RIPS works with d_{ABCD} range measures and not pairwise distances. What are the possible values d_{ABCD} can take? Since all four terms in the equation are between 0 and the maximum interferometric radio range (r), it is easy to see that

$$-2r \leq d_{ABCD} \leq 2r. \quad (19)$$

7.2 Experimental setup

Due to the difficulties involved in establishing the ground truth very accurately, so far we have only experimented with smaller scale setups. The data we analyze in the following sections were gathered using 16 nodes in a 4x4 grid deployed in a flat grassy area with no obstructions. The overall area covered was 18x18 meters. In order to localize nodes in 2D we selected 3 anchor points. They were picked randomly while making sure that they did not fall in a line.

The grid was selected because of the relative ease of setting up the positions precisely. While measuring the edges with a tape measure can be done well enough, keeping all the angles right is harder. We estimate that the accuracy of node placement is within 5cm. Note that we do not use the fact that the nodes are arranged in a grid at any point in our positioning method.

7.3 Frequency and phase accuracy

The performance of our frequency and phase detection algorithm is comparable to a high resolution (1 Hz) DFT-based approach. The justification of DFT-based tone parameter estimation and its relationship to the maximum likelihood estimator can be found in [13].

Figure 6 shows frequency and relative phase offset results of frequency tuning observed on a pair of motes (one of the senders changed its carrier frequency in small increments). The number of RSSI samples (measurement interval) limits both approaches at low frequencies. In the “normal operating range,” however, the mote implementation performs surprisingly well. On the frequency diagram both methods

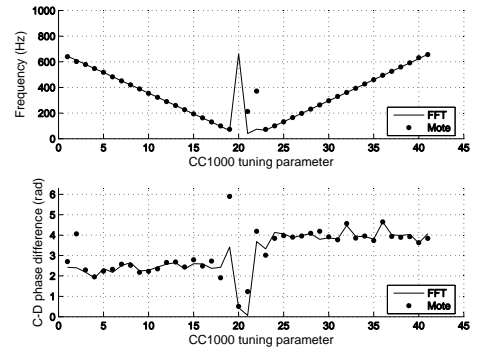


Figure 6: Comparison of frequency and phase results and high resolution DFT estimation at different interference frequencies.

closely reveal the ideal “v-shaped” curve. Phase difference measurements have significantly more noise (the ideal response would be a constant value flipped at 0 Hz), as we expected in Section 4.4.

The average error of relative phase offset measurement is illustrated in Figure 7. In the experimental setup we fixed one pair of motes as transmitters and performed frequency tuning around 0 Hz interference. We repeated the same experiment 30 times. For all $\binom{14}{2}$ pairs we calculated the median phase and the average deviation at each frequency. Next, we calculated the average value of these deviations. By using the amplitude value as a quality indicator of measurements, the average error can be drastically reduced.




Figure 7: Mean deviation of phase measurements using different filtering thresholds.

7.4 Ranging accuracy

The error distribution of the calculated d_{ABCD} ranges can be approximated by the superposition of a set of Gaussian distributions with centers at integer multiples of the wavelength of the carrier frequency. Depending on the experimental setup, about half of the range estimates have less than one quarter wavelength error with the remaining values shifted by one or two wavelengths as shown in Figure 8. The ratio of the good and shifted values can be significantly improved using simple filtering techniques while keeping the number of range estimates high enough to enable accurate localization of the nodes.

The algorithm measuring the interference signal frequency and phase also determines the average amplitude of the signal. The amplitude shows strong correlation with the error of the range estimate. Currently, we use a constant amplitude threshold (12% of maximum A/D range) to discard measurements with low SNR. This filtering can be carried out locally on each mote acting as receiver, but currently it is done on the base station.

The interference signal is measured by all the nodes in radio range. Due to measurement errors, the frequency estimates will vary at different motes. Nodes that measure the frequency with a large error will likely have a bad phase estimate also. Therefore, these measurements need to be filtered out as well. The filtering is carried out by identifying a narrow frequency window that has the maximum number of frequency estimates in it. All the measurements outside of this window are discarded. This process can also be carried out on the motes, but it would require communication among the active receivers. Currently it is done on the base station.

Finally, we calculate the range for a given pair of transmitter receiver pairs only if the number of frequency channels with good phase offset measurements is higher than a threshold. Currently this limit is set to 10.

After these three filtering stages, the ratio of the measurements with less than one quarter wavelength error can be improved by approximately 50% as illustrated on Figure 9.




Figure 8: Error distribution of all the ranges.




Figure 9: Error distribution of the filtered ranges.

Figure 10 shows the central portion of the error distribution after filtering. The accuracy of these over 2000 measurements clearly demonstrates the potentially extreme high-

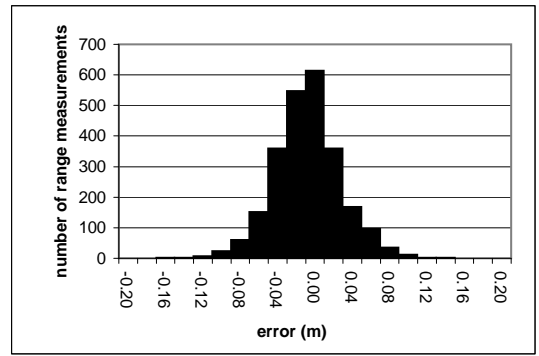


Figure 10: Central portion of the error distribution of the filtered ranges.

precision of overall localization using RIPS if one can eliminate the “side lobes” of the distribution. We can either develop more advanced filtering methods to discard measurements with full wavelength errors, increase the accuracy of the phase measurements, or increase the frequency band beyond the [400, 460] MHz range. Even a small improvement in phase estimation accuracy could potentially dramatically increase the ratio of good to bad measurements. Intuitively there is a threshold in the phase measurement error where it is not large enough to cause the range estimator to miss by a full wavelength. An analytical evaluation is needed to quantify this relationship.

7.5 Localization accuracy

We ran a localization experiment using the setup described above utilizing the filtered ranging data shown in Figure 9. The genetic optimization procedure ran for 2 minutes. The error distribution of the resulting localization is shown in Figure 11. The average accuracy was 3 cm, while the largest error was approximately 6 cm. The results are shown in Figure 12 with the three anchor nodes depicted by large circles. At this resolution and localization accuracy the small circles showing the actual and estimated positions overlap.




Figure 11: Error distribution of localization.

To test the scalability of the approach we reran the localization utilizing only 20% of the raw ranging data. We selected 48 transmitter pairs out of the possible $\binom{16}{2} = 240$ randomly. After filtering, approximately 1000 measurements remained with 28% of them shifted by integer multiples of

the wavelengths. This ratio is about the same as for the whole data set. The localization algorithm achieved 5 cm average accuracy, while the worst error remained under 10 cm.




Figure 12: Localization results.

Note that the estimated accuracy of the ground truth and the accuracy of the localization results are comparable. Therefore, these numbers are not exact; they are just indications of the very high precision RIPS can achieve.

7.6 Latency

There are $\binom{16}{2}\binom{14}{2}$ different configurations of transmitter/receiver pairs in the 16-node experiment described above. We use three different amplitude combinations for any given four-node setup. Therefore, there are altogether approximately 32,000 measurements carried out. Note that not all of these measurements are independent but we wanted to gather as much ranging data as possible. There is quite a bit of concurrency, as for any particular transmitter pair, all receivers perform their measurements in parallel. On the other hand, each transmitter pair needs to do multiple measurements for frequency tuning and then multiple frequency channels are used for the actual measurements. The tuning algorithm, range calculations and localization are carried out on the base station, so there is a large amount of data that is shipped to the root of the network.

Currently, the entire process takes about 80 minutes. This can be sped up significantly. Localization does not need this amount of data; potentially an order of magnitude less would work well enough. If we use only one fifth of the possible transmitter pairs—to be conservative—that would immediately decrease the time to below twenty minutes. Furthermore, the tuning algorithm can be implemented on the motes. It needs to run on one of the receivers of an interference signal. The results then need to be communicated back to the transmitters. This would help scalability in larger setups since the base station need not be involved in tuning at all and hence, message routes would be much shorter. As mentioned earlier, data filtering based on the interference signal amplitude can be done locally on the mote. Additional filtering requires communication between the different receivers of the same interference signal. In a small setup, it would not provide any speedup. For a large deployment, however, this would significantly increase the scalability.

Additional optimization could involve decreasing the number of frequency tuning steps to the minimum necessary.

The number of channels used for the actual phase measurements could also be decreased somewhat if the phase estimation accuracy can be improved. Finally, the mote implementation code is not fully optimized in this first version of RIPS. We estimate that the entire localization process could be carried out in under 5 minutes for smaller scale setups. Large setups would require more time, but the process scales well because the network can be automatically divided into sets of non-overlapping regions determined by the radio interferometric range where the procedure can be carried out concurrently. The design of an efficient scheduling algorithm is an area of further research.

8. CONCLUSIONS AND FUTURE WORK

The Radio Interferometric Position System (RIPS) represents a significant advance in node localization in WSNs. It achieves high accuracy and long range simultaneously, supports 3D localization and does not require extra hardware or calibration. The key enabling ideas behind this performance are the application of two transmitters to create an interference signal, the measurements of the relative phase offset at two receivers cancelling out many sources of error, and the fact that we measure the phase of a low-frequency signal, yet it relates to the wavelength of the high-frequency carrier signal.

Our prototype implementation was enabled by the highly configurable Chipcon CC1000 radio. Hence, it is available on the MICA2 and XSM platforms, but not on MICAZ-s or Telos-es. In the future we foresee the implementation of the ranging technique in hardware: the radio frequency should be tunable and the radio chip needs to support frequency and phase measurements natively. This would simplify the localization service, speed up the operation and provide higher accuracy.

Our current work involves the evaluation of the technique indoors. We are researching how RF multipath effects distort the measurements. Efficient scheduling of the individual measurements in order to minimize localization time in a large scale WSN deployment is another area of study. Finally, we plan to verify the performance of RIPS in a large scale 3D field experiment in the near future.

9. ACKNOWLEDGMENT

The authors would like to thank Janos Sallai, Manish Kushwaha, Martin Haenggi, Vijay Raghavan, Janos Sztiapanovits and Lambert Meertens for their valuable contribution to this work. We also thank our shepherd, Andreas Savvides for his guidance and the anonymous reviewers for their constructive criticism.

The DARPA/IXO NEST program (F33615-01-C-1903) has supported the research described in this paper.

10. REFERENCES

- [1] Gy. Simon, M. Maróti, Á. Lédeczi et al., *Sensor network-based countersniper system*, in Proc. of ACM 2nd Conference on Embedded Networked Sensor Systems (SenSys), November 2004, 1–12.
- [2] T. He, S. Krishnamurthy, J. Stankovic, et al. *An Energy-Efficient Surveillance System Using Wireless Sensor Networks*, in Proc. of MobiSys'04, June 2004.

- [3] H. Wang, J. Elson, L. Girod, D. Estrin, K. Yao, *Target classification and localization in a habitat monitoring application*, In Proc. of the IEEE ICASSP, April 2003.
- [4] Y. Kwon, K. Mechitov, S. Sundresh, W. Kim and G. Agha, *Resilient localization for sensor networks in outdoor environments*, Technical Report UIUCDCS-R-2004-2449, Department of Computer Science, University of Illinois at Urbana Champaign, 2004.
- [5] K. Whitehouse and D. Culler, *Calibration as parameter estimation in sensor networks*, In Proc. of the ACM international workshop on wireless sensor networks and applications, Atlanta, GA, 2002.
- [6] J. Sallai, Gy. Balogh, M. Maróti, Á. Lédeczi and B. Kusý, *Acoustic ranging in resource-constrained sensor networks*, in Proc. of International Conference on Wireless and Mobile Computing (ICWN), June 2004.
- [7] N.B. Priyantha, A. Chakraborty and H. Balakrishnan, *The Cricket Location-Support System*, In Proc. of MobiCom 2000: The Sixth Annual International Conference on Mobile Computing and Networking, Boston, MA, August 2000.
- [8] K. Whitehouse, C. Karlof, A. Woo, F. Jiang and D. Culler, *The Effects of Ranging Noise on Multihop Localization: An Empirical Study*, in Proc. of IPSN, Los Angeles, CA, April, 2005.
- [9] J. Hightower, R. Want and G. Borriello, *SpotON: An Indoor 3D Location Sensing Technology Based on RF Signal Strength*, University of Washington, Tech. Rep. UW CSE 2000-02-02, Feb 2000.
- [10] K. Yedavalli, B. Krishnamachari, S. Ravula and B. Srinivasan: *Ecolocation: A Technique for RF Based Localization in Wireless Sensor Networks*, in Proc. of IPSN, Los Angeles, CA, April, 2005.
- [11] D. Niculescu and B. Nath, *Ad Hoc Positioning System (APS)*, in Proc. of IEEE GLOBECOM 2001, San Antonio, Nov 2001.
- [12] L. Meertens, *The Dimension of the Vector Space Spanned by Sets of Radio-Interferometric Measurements*, Technical Report KES.U.05.02, Kestrel Institute, July 2005.
- [13] D. Rife, R. Boorstyn, *Single tone parameter estimation from discrete-time observations*, IEEE Transactions on Information Theory, Volume 20, Issue 5, pp. 591–598, September 1974.
- [14] S.A. Tretter, *Estimating the Frequency of a Noisy Sinusoid by Linear Regression*, IEEE Transactions on Information Theory, Volume 31, Issue 6, pp. 832–835, November 1985.
- [15] M. Maróti, B. Kusý, Gy. Simon and Á. Lédeczi, *The flooding time synchronization protocol*, in Proc. of ACM 2nd Conference on Embedded Networked Sensor Systems (SenSys), November 2004, 39–49.
- [16] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister, *System architecture directions for networked sensors*, in Proc. of ASPLOS 2000, Cambridge, MA, November 2000.
- [17] J. Polastre, R. Szewczyk, C. Sharp, D. Culler, *The Mote Revolution: Low Power Wireless Sensor Network Devices*, in Proc. of Hot Chips 16, August 2004.
- [18] Chipcon Inc., *Chipcon CC1000 Data Sheet*, ver. 2.2, <http://www.chipcon.com/>.