Stochastic Radio Interferometric Positioning in the 2.4 GHz Range

B.J.Dil
Pervasive Systems Group
University of Twente
PO-Box 217, 7500 AE Enschede, The Netherlands
B.J.Dil@utwente.nl

P.J.M.Havinga
Pervasive Systems Group
University of Twente
PO-Box 217, 7500 AE Enschede, The Netherlands
P.J.M.Havinga@utwente.nl

Abstract

This paper presents a novel Radio Interferometric Positioning System (RIPS), which we call Stochastic RIPS (SRIPS). Although RIPS provides centimeter accuracy, it is still not widely adopted due to (1) the limited set of suitable radio platforms and (2) the relatively long measurement and calibration times. SRIPS overcomes these practical limitations by (1) omitting the calibration phase of the existing RIPS and by (2) applying a novel positioning algorithm. SRIPS exploits the phenomenon of the small but stable difference between two transmitted frequencies that often exists when two radios are tuned to the same frequency. We obtain an experimental measure for this stability. This approach enables the implementation of RIPS on commonly available radio platforms, such as the CC2430, because fine-tuning in small steps relative to the beat frequencies for calibration is not required. In addition, we show that SRIPS calculates the position that provides the best fit to the set of measurements, given the underlying statistical and propagation models. Therefore, SRIPS converges more accurately to the true locations in a variety of situations of practical interest. Experiments in a $20 \times 20m^2$ set-up verify this and show that our SRIPS CC2430 implementation reduces the number of required measurements by a factor of three, and it reduces the measurement time to less than 0.1 seconds, while providing accuracy similar to that of the existing RIPS implementation on the CC1000 platform, which requires seconds.

Categories and Subject Descriptors

C.2.4 [**Computer-Communications Networks**]: Distributed Applications

General Terms

Algorithms, Experimentation, Theory

Keywords

Wireless Sensor Networks, Radio Interferometric Localization, RSS-based Localization, Tracking, Localization

1 Introduction

This paper describes a novel approach for Radio Interferometric Positioning of devices in wireless networks that we call Stochastic RIPS, or SRIPS. Localization in wireless networks is the process of finding a physical location in an automated manner using wireless communication.

Practically, localization can be a stand-alone application (e.g. inventory tracking in a distribution center), or it can provide support to the network service (e.g. routing). Today, such applications have evolved into real-time location systems (RTLS) using a wide range of wireless technologies. Many of these localization applications are based on Received Signal Strength (RSS) measurements because RSS information is obtained without additional hardware and energy costs. Other localization systems use techniques such as Time Difference Of Arrival (TDOA), Time Of Flight (TOF), and Angle Of Arrival (AOA). In general, these techniques can be more accurate than RSS-based localization, but they require specialized hardware, more processing, more communication, and thus more energy (e.g. [6]).

Radio Interferometric Positioning does not depend on the received signal strength as a measure for the location of unknown nodes as in standard RSS-based localization. Rather, it depends on the stability of two slightly different frequencies of a pair of carrier waves and on the nonlinear permittivity of a pair of receiving antennas to generate beat signals. Electromagnetic Interferometric Positioning was first developed with the same propagation models in the optical regime with the advent of Zeeman Lasers ([1]). The Zeeman Laser Interferometer generates two laser frequencies relatively close to each other, generating a frequency beat at the detector. The stability of the frequency beat is a direct measure for the positioning accuracy ([1]). As this frequency beat is derived from the same source, you can obtain nanometer accuracy over a range of one meter ([3]).

RIPS ([8]) relies on two independent sources. These sources, called sender pairs, simultaneously transmit unmodulated carrier waves at slightly different frequencies that must be stable during the measurement time of one frequency beat signal. Receiving node pairs that are within transmission range measure the energy of the frequency beat

SenSys'11, November 1–4, 2011, Seattle, WA, USA.

signal. The phase offsets of the beat signals are a function of the distances between the nodes involved relative to the carrier wavelengths. These phase offset measurements are performed for all sending and receiving node pairs over a range of discrete carrier frequencies. This set of phase offset measurements serves to calculate the unique geometrical path differences between each sending and receiving node pair. Each path difference is called a q-range ([9]). RIPS calculates the position estimate by fitting it to the q-ranges. [8] shows that RIPS can achieve centimeter accuracy over a range of 100 meters.

[8] calculates each q-range separately, by fitting the phase offset measurements of each sending and receiving pair separately. Then it fits the position estimate to these calculated q-ranges. This approach does not always provide the best fit given the set of phase offset measurements and underlying statistical and propagation models. A more rigorous approach is to fit all phase offset measurements together, rather than fitting the phase offset measurements of each sending and receiving pair separately. That approach should converge to the best fit of all phase offset measurements, and that is what SRIPS does. It fits all the phase offset measurements by evaluating the q-ranges as stochastic functions instead of deterministic ones.

RIPS achieves its centimeter accuracy at the cost of (1) a strict requirement on the radio platform that limits the set of suitable radio platforms and (2) relatively long measuring and calibration times. SRIPS overcomes these practical limitations by (1) omitting the calibration phase and by (2) applying its new algorithm. The RIPS calibration phase requires that the radio tune its frequency in steps smaller than the desired frequency beat, so that the frequency beat can be calibrated. Most Commercial-Off-The- Shelf (COTS) radios do not comply with this tuning requirement. As a result, RIPS can be implemented only on a few radio platforms, such as the CC1000 ([8]-[16]). To our knowledge, there are no IEEE 802.15.4 radios in the 2.4 GHz range that comply with the frequency tuning requirement of the original RIPS implementation (e.g. [21]-[24]).

One of the reasons we implemented SRIPS on a CC2430 radio platform is that it operates in the 2.4 GHz range. The CC1000 platform ([20]) operates in the 400 and 800/900 MHz range. The allocated bandwidth (ISM band) within these frequencies is non-overlapping for different regions in the world. For instance, in Europe the allocated bandwidths are limited to 433.05-434.79 MHz and $868.0 - 868.6$ MHz. The existing RIPS implementation uses a bandwidth of 60 MHz ($[400, 460]$ MHz, [8]), which is outside the allocated bandwidth below the 1 GHz range.

The contributions of this paper are as follows:

- The accuracy of Electromagnetic Interferometric Positioning depends on the stability of the generated frequency beat ([1]). This paper experimentally verifies that the stability of the generated frequency beat signal is sufficient for at least 80 milliseconds, which provides the required stability for a position accuracy of centimeters.
- We introduce a novel Radio Interferometric Positioning System called SRIPS. Measurements on a CC2430 platform in a 20×20 m^2 outdoor environment show that SRIPS provides an accuracy of ~ 0.3 meters. We show that our SRIPS CC2430 implementation provides results similar to those of the existing RIPS CC1000 implementation, while reducing the measurement time from 1 to 0.06 seconds and reducing the number of measurements by a factor of 3.
- We compare the performance of our SRIPS algorithm with a typical RSS-based localization algorithm ([6]). Measurements on a CC2430 platform in a 20×20 m^2 outdoor environment show that SRIPS improves the accuracy from ~ 6 to ~ 0.3 meters compared with a typical RSS-based positioning algorithm .

Section 2 of this paper reviews the existing work in the field of Optical and Radio Interferometric Localization. Section 3 describes how we implemented the Radio Interferometric Position System on the CC2430 platform. In addition, it shows that the CC2430 platform provides the required stability for a position accuracy of centimeters. Section 4 describes the SRIPS positioning algorithm and shows how our algorithm copes with the varying distance estimate accuracies. Section 5 evaluates RIPS and SRIPS on a CC2430 platform. In addition, it compares our CC2430 SRIPS implementation with the existing CC1000 RIPS implementation. Section 6 provides a conclusion.

2 Background

This section presents a brief summary on Interferometric Positioning. Because the propagation models for Optical and Radio Interferometric Positioning are essentially the same, and because the former was developed thirty years earlier, we start with positioning based on Optical Interferometry. Then we describe the measurement set-up we use throughout this paper. Section 2.3 provides the theoretical background on Radio Interferometric Positioning. At the end of this section, we compare the process flows of RIPS and SRIPS.

2.1 Interferometric Positioning in the Optical Regime

In Optical Interferometric Positioning Systems, the two interference signals are generated by one coherent laser source that is split into two Zeeman modes ([1]). These two modes are split by a special beam splitter into a reference and signal mode, with the signal mode reflected by a moving target that must be positioned. The frequency beat (Δf) of these two modes is on the order of a few hundred MHz. The positioning is derived from the Doppler shifted phase difference between the two modes. The resolving power of this positioning device is directly related to the stability of the beat signal: $\frac{\delta(\Delta f)}{\Delta f}$ ([1]). Hence, the error increases linearly with the distance: $(d_{BC} - d_{BD}) \cdot \frac{\delta(\Delta f)}{\Delta f}$ ([1]). The Zeeman Laser provides a high frequency beat stability because both modes are derived from the same source. Inherent instabilities in the source are canceled out at the detector. This results in nanometer resolving power on a range of one meter ([3]). RIPS uses two independent radio transmitters with two independent receiving pairs, and this seriously limits the stability of the resulting beat signals compared to the Optical Interferometer. As we see later, the resulting resolving power is on the order of $10^{-2} \ldots 10^{-3}$ as compared with 10^{-9} in the optical case. In this paper, we empirically measure the resolving power of our radio platform by determining the stability of the beat signals (see Section 3.2).

2.2 Radio Interferometric Positioning Set-up

This paper focuses on Radio Interferometric Positioning of a target node using a network infrastructure. This set-up distinguishes two types of nodes: the infrastructure nodes and the target nodes. The infrastructure nodes know their location and support the target nodes to position themselves. In general, this set-up provides decimeter accuracy in a measurement time of one second (see Section 5.5). The target node can be either a transmitter or receiver. We implement the target-as-receiver implementation, so that we can position multiple target nodes in parallel without increasing the measurement time. This scalability is not possible with the target-as-sender implementation, because in that set-up the measurement time increases with the number of target nodes ([11] and [12]). We consider the scalability of increasing the number of target nodes without affecting the measurement time as an important aspect in wireless positioning systems.

Figure 1 shows the RIPS/SRIPS set-up used throughout this paper. The circles represent the infrastructure nodes (A/B/C/E) and the triangle (D) represents the target node. Red indicates transmitting nodes and green receiving nodes. The diagram in Figure 2 shows which measurements are performed to estimate the position of one or several target nodes:

- **Phase Offset Measurements**
  Two infrastructure nodes form a sender pair and generate a beat signal; two or more receivers sample the beat signal with 'O' sample points. RIPS uses 256 sample points. We vary this amount for analysis (see Section 5). The CC1000 has a sampling frequency of 9 kHz, and the CC2430 has a sampling frequency of 62.5 kHz. The beat frequency and phase offset is then calculated from these sampled beat signals for each receiver pair. Each receiver pair includes one infrastructure node and one target node. Note that all participating nodes must be accurately synchronized (microsecond synchronization, [8]).
- **Measurement round**
  This phase consists of performing N phase offset measurements at N different frequencies to calculate the q-range between the sending and receiving node pair (see Section 2.3). However, calculating q-ranges is computationally intensive. Therefore, the calculated phases and frequencies are sent to a computer for post-processing at the end of each measurement round ([8]-[13]). Each measurement round is identified by a unique sender pair.
- **Positioning of the target node**
  The position of the target node is calculated on a central computer on the basis of the data of two or more measurement rounds (independent q-ranges).

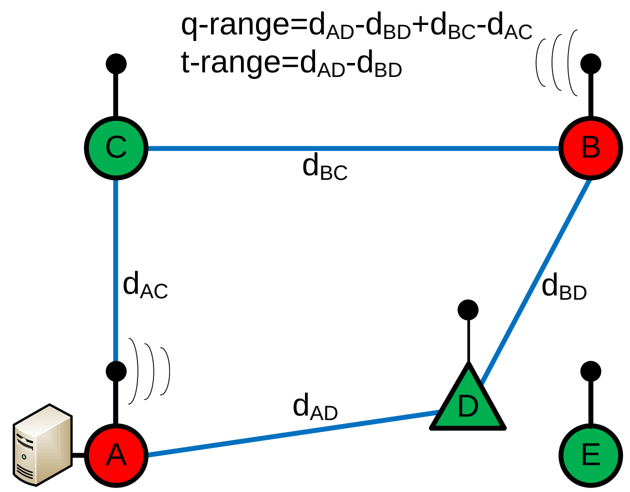

Figure 1. Measurement set-up: red = transmitting, green = receiving, circle = infrastructure node, triangle = target node

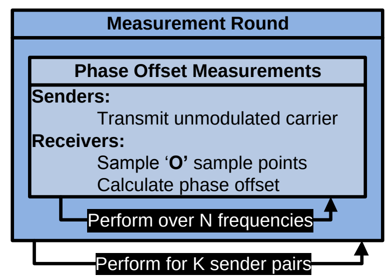

Measurement Round (K)	Senders	Receivers	Frequencies (N)
1	A / B	C / E / D	38
2	A / C	B / E / D	38
3	A / E	B / C / D	38
4	B / C	A / E / D	38
5	B / E	A / C / D	38
6	C / E	A / B / D	38

Figure 2. Measurement phases

The table in Figure 2 shows which nodes are sending/receiving during each measurement round in the measurement set-up shown in Figure 1. For instance in the first measurement round, infrastructure nodes A and B are transmitting unmodulated carriers. Nodes C, E and D measure the beat signal ($R = 3$ Receivers). Each beat signal is measured over 38 frequencies in our set-up. Figure 2 shows that there are $K = 6$ measurement rounds, each allowing for two receiver pairs. In this example, we assume that the receivers sample the beat signal with $O = 500$ sample points. Then the total number of sample points used to position the target node equals:

$$K \cdot R \cdot N \cdot O = 6 \cdot 3 \cdot 38 \cdot 500 = 342000 \tag{1}$$

The total measurement time to position the target node does not depend on the number of receivers:

$$\text{Total measurement time} = K \cdot N \cdot O \cdot \frac{1}{f_s} \tag{2}$$

Here f_s is the sampling frequency of the hardware used. Therefore, increasing the number of receiving infrastructure nodes increases the total number of sample points used to position the target node. This increases the positioning accuracy (e.g. [12]), without increasing the total measurement time, and therefore energy consumption of the target nodes.

2.3 Radio Interferometric Positioning

Nodes A and B in Figure 1 transmit an unmodulated carrier signal with frequencies f_A and f_B. Receivers C and D measure the energy of the composite signal with a frequency beat of $\Delta f = |f_A - f_B|$. The measured phase offset is a function of the distances between sender pair A/B and receiver pair C/D, assuming that $f_A > f_B$:

$$\Delta\varphi_i = 2\pi\left(\frac{d_{AD}-d_{AC}}{\lambda_A} - \frac{d_{BD}-d_{BC}}{\lambda_B}\right) \mod (2\pi) \tag{3}$$

$$\approx 2\pi\left(\frac{d_{ABCD}}{\lambda_i}\right) \mod (2\pi)\ . \tag{4}$$

Here:

$$\text{q-range} = d_{ABCD} = d_{AD} - d_{BD} + d_{BC} - d_{AC}\ , \tag{5}$$

and:

$$\text{t-range} = d_{ABCD} - d_{BC} + d_{AC} = d_{AD} - d_{BD}\ . \tag{6}$$

Here $\Delta\varphi_i$ is the relative phase offset; d_{AD} is the distance between node A and D; [8] defines λ_i as $\lambda_i = \frac{2c}{f_A+f_B}$ with c representing the speed of light. We define d_{ABCD} as the q-range, as in [9]. The locations of infrastructure nodes A, B and C are known, so the distances between these nodes are also known (d_{BC} and d_{AC}). We use this information to transform the q-range into the t-range, so that all the values of the variables on the left side of the equation are known. Throughout this paper, we use the q-range in the equations and the t-range in the figures, because the t-range is independent of the location of the receiving infrastructure node. This is useful for comparing q-ranges of the same sender pair, as we show in Section 4.3.

Equation 3 does not define a unique solution for the q-range (d_{ABCD}) due to $\mod (2\pi)$-related ambiguity of the q-range. Therefore, RIPS implementations perform phase offset measurements over N frequencies for the same sender pair: $f_1 \ldots f_N$. The squared error of a q-range estimate is given by ([9]):

$$\text{ERROR}(\text{q-range}_j) = \sum_{i=1}^{N} (\text{q-range}_j - d_{i,j})^2 \tag{7}$$

Where:

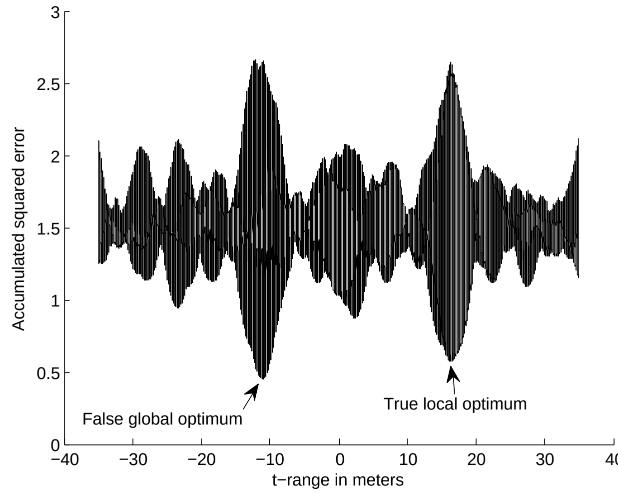

Figure 3. t-range error distribution, local optimum

$$d_{i,j} = \text{round}\left(\frac{\text{q-range}_j - \gamma_{i,j}}{\lambda_{i,j}}\right) \cdot \lambda_{i,j} + \gamma_{i,j}$$

minimizes Equation 7 for a given q-range; $\gamma_{i,j}$ represents the phase offset relative to the wavelength $\gamma_{i,j} = \lambda_{i,j}\frac{\Delta\varphi_{i,j}}{2\pi}$. Figure 3 shows the t-range error distribution calculated by Equation 7 as a function of the t-range.

The Most-Likelihood-Estimator of the q-range is the global minimum of the q-range error distribution, minimizing the squared error between the calculated and measured phase offsets:

$$\begin{aligned} q_{\text{est},j} &= \arg \min_{\text{q-range}_j} \sum_{i=1}^{N} (\Delta\varphi_{i,j} - \widehat{\Delta\varphi_{i,j}})^2 \\ &= \arg \min_{\text{q-range}_j} \sum_{i=1}^{N} (\text{q-range}_j - d_{i,j})^2 \\ &= \arg \min_{\text{q-range}_j} \text{ERROR}(\text{q-range}_j) \end{aligned} \tag{8}$$

Most RIPS algorithms use this deterministic approach for estimating one q-range on the basis of the calculated q-range distribution (e.g. [8]). Figure 3 shows an example of a t-range error distribution of one sending pair. It shows that the real t-range is located in a local minimum instead of the global minimum. Such error distributions seriously limit the accuracy of RIPS but can be handled by SRIPS. We discuss this further in Section 4.

The position is estimated by minimizing the squared difference between the calculated and estimated q-ranges:

$$\{\hat{x}, \hat{y}\} = \arg \min_{\hat{x},\hat{y}} \sum_{j=1}^{M} (q_{\text{est},j} - \widehat{\text{q-range}}_j)^2 \tag{9}$$

Here, M is the number of q-range estimates; $\widehat{\text{q-range}}_j$ is the calculated q-range as a function of the estimated position $(\hat{x}, \hat{y})$. Although the present literature shows that Equation

9 converges with increasing N (phase offset measurements over different frequencies) and M (q-range estimates), they don't necessarily converge to the true positions, as shown in Section 4.

2.4 Process flow of RIPS and SRIPS

Figure 4 shows the individual phases of the CC1000 RIPS and the CC2430 SRIPS implementations:

- **Calibration Phase**
  The calibration phase ensures that the frequency beat is measurable given the user-defined sample time (28...40 ms) and the hardware-defined sampling rate (9 KHz). However, the calibration phase requires that the radio set the frequency in small steps relative to the desired beat frequencies. The CC1000 can tune its frequency in steps of 65 Hertz, which is sufficient for a frequency beat of 300 to 400 Hertz ([8]). The CC2430 does not comply with this tuning requirement; therefore, SRIPS follows another approach (see Section 3.1).
- **Measurement Phase**
  In this phase, the receivers measure the phase and frequency of the beat signals. RIPS and SRIPS use a simple threshold-crossing technique ([2]) for estimating the phase offset between the receivers, to keep the computational costs low.
- **Calculate Distance Distributions**
  The q-range error distribution is calculated using Equation 7.
- **Estimate Distances (Deterministic)**
  In this phase, the q-range is estimated on the basis of the calculated q-range error distribution. Figure 3 shows an example where the real q-range lies in a local minimum, which is a known problem of RIPS (e.g. [9]). Most RIPS algorithms assume that the real q-range is determined by this global minimum (Equation 8). [9] and [11] try to solve this problem, however they are dependent on parameter settings that can only be determined empirically. SRIPS does not estimate q-ranges, but evaluates the q-range error distribution directly for estimating the position. Therefore, it does not have this distance estimation phase (for details we refer to Section 4).
- **RIPS and SRIPS Positioning Algorithms**
  In this phase, RIPS estimates the position on the basis of the q-range estimates (Equation 9). SRIPS estimates the position on the basis of the calculated q-range error distributions.

Figure 4 shows that the only difference between RIPS and SRIPS is the omitted calibration phase and the localization algorithm. Otherwise, RIPS and SRIPS perform the same measurements and use the same input for localization.

3 SRIPS Measurement Phase and Error Characterization

In this section, we focus on the measurement phase. In addition, we empirically verify whether the CC2430 hardware is suitable for radio interferometric positioning.

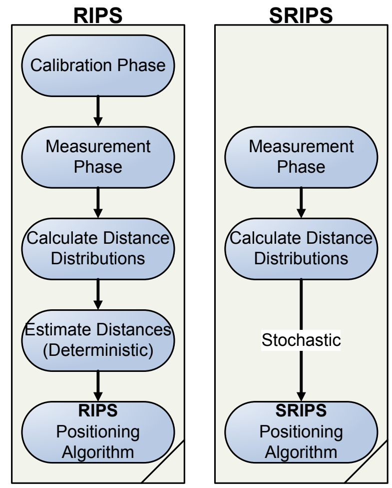

Figure 4. RIPS vs SRIPS

3.1 CC2430 Measurement Phase

A typical target-as-receiver RIPS and SRIPS set-up is shown in Figure 1. Although our SRIPS implementation tunes each sender pair to the same frequency, in practice these frequencies are never exactly the same. There is often a small but stable difference between the transmitted frequencies caused by small and stable differences between the crystal oscillators. We measure the beat frequencies at receiver pairs in our user-defined measurement time of 8 ms per frequency at the hardware-defined sampling rate of 62.5 KHz.

The measurement time per frequency determines the lower bound of the measurable frequency beat: $\frac{1}{0.008} \approx 125$ Hz. The Nyquist frequency determines the upper bound of the measurable frequency beat: $\frac{62.5}{2} \approx 30$ KHz. Measurements with 48 different sender pairs show that 85% of these pairs generate measurable frequency beat signals.

Our CC2430 SRIPS test bed contains a master node, which controls and collects the SRIPS measurements. The master node can be a transmitter or receiver. Our SRIPS implementation consists of the following steps in order to perform one phase offset measurement:

1. The master node sends a synchronization message, which synchronizes the nodes and identifies the transmitters and receivers. We use the CC2430 MAC controller in order to obtain at least 208 ns synchronization accuracy at the reception of the synchronization message ([21]). The time (1.4 ms) between the synchronization message and sampling the beat signal introduces an additional error. This error depends on the accuracy of the crystal oscillator and the 1.4 ms waiting time. [17] shows that this introduces an extra error of approximately 60 ns.
2. 1 ms after receiving the synchronization message, the transmitters start sending unmodulated carrier signals.
3. 1.4 ms after receiving the synchronization message, the receivers start sampling the beat signal. We use the CC2430 DMA controller to perform the measurements

Platform	CC1000	CC2430
Calibration	Yes	Not available
Sample time	28.4…40 ms	0.8…8 ms
RSS sampling rate	9 kHz	62.5 kHz
Frequency beat	200…800 Hz	0.2…14 kHz
Frequency range	400 - 800/900 MHz	2.4 GHz
Positioning algorithm	RIPS	SRIPS

Table 1. Platform characteristics

with clock tick accuracy. This means that the accuracy of the crystal oscillator determines the measurement jitter. [17] shows that the maximum measurement jitter is on the order of 320 ns when sampling 500 consecutive sample points over 8 ms.

4. At the end of each phase offset measurement, the master node collects the 500 sample points from each receiver for analysis.

RIPS and SRIPS use the same methodology to obtain the required phase offset measurements. The only difference is that in our current SRIPS implementation, we send the raw sample points to a computer for analysis. In a commercial implementation, these computations can be distributed to the local processors of the nodes. Table 1 compares the characteristics of the CC1000 and CC2430 platform.

3.2 CC2430 Error Characterization

Before we further analyze SRIPS and its performance in a practical localization set-up, we perform several frequency beat measurements in a limited set-up to show that our CC2430 hardware with a 20 ppm crystal oscillator is suitable for radio interferometric positioning. This limited set-up consists of four nodes placed in the corners of a 1×1 m^2 square. The short-range and line-of-sight measurements minimize the influence of the environment on the received signals, so that we practically measure the performance of the CC2430 hardware. We place all radios on tripods at the same height of 1.5 meters, and we do not place objects in the vicinity of the radios, to minimize the influence of interfering reflections ([9]). The conditions during the measurements are static (temperature, humidity, no moving objects). Receiving nodes sample the beat signals with 5000 sample points at a sample rate of 62.5 KHz. All nodes are tuned over 38 frequencies in a bandwidth of 2.406…2.480 GHz. The frequency beat signals are measured with all possible sender and receiver pairs (6). Each node is equipped with a widely used omnidirectional dipole antenna with a vertical orientation.

Minimum bandwidth is the inverse of the maximum coherence time or stability of the beat signals. Therefore, increasing the measurement time beyond the coherence time does not provide reliable phase measurements as expressed by Equation 3. We calculate the bandwidth of the first 500 sample points and over all 5000 sample points. Figures 5 and 6 show the Fourier transform of the corresponding frequency beat signals. In each figure, the red curve represents the Fourier transform of the first 500 sample points and the black curve represents the Fourier transform of the 5000 sample points. These figures clearly show that the bandwidth of the spectrum of the black curve ($\sim$ 20 Hz) is a factor of 10 smaller than the corresponding bandwidth of the red curve ($\sim$ 200 Hz). Only coherent signals reduce the frequency bandwidth proportional to the measurement time. This is a direct result of the Heisenberg-Pauli-Weyl inequality, because the product of the variance of time and frequency remains constant. This means that the stability of our beat signals is at least 80 ms.

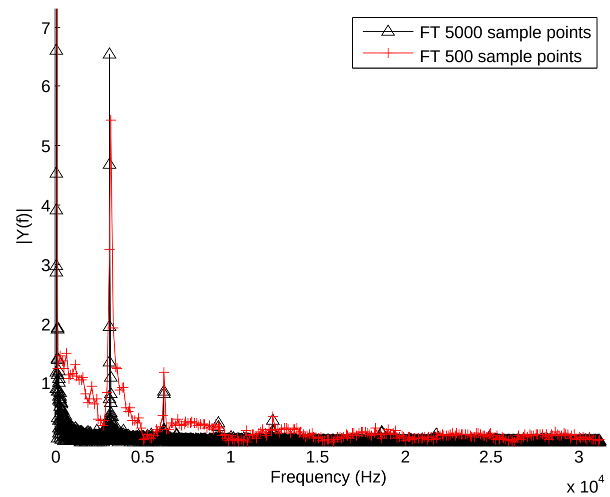

Figure 5. Single-Sided Amplitude Spectrum, FT stands for Fourier Transform

[8] states that the interferometric range of the CC1000 carrier waves is on the order of a few hundred meters. Our SRIPS algorithm measures a frequency stability of at least 80 ms, which corresponds to a theoretical coherence length of at least 2400 km for the plane waves of our propagation model that have an infinite range in free space. However, the energy transmitted from the antenna is not concentrated into a narrow beam, so that in the far field this energy decreases inversely proportionally to the square of the distance from the radio transmitters. This path loss in free space limits the application of RIPS and SRIPS to localization areas with a diagonal of a few hundred meters. Therefore, both RIPS and SRIPS can determine the positions with a relative accuracy of roughly $(d_{BC} - d_{BD}) \cdot \frac{\delta(\Delta f)}{\Delta f} \approx (d_{BC} - d_{BD}) \cdot 10^{-3}$ when the measurement times are adapted to approach the coherence times of the radio transmitters. In principle, it is possible to

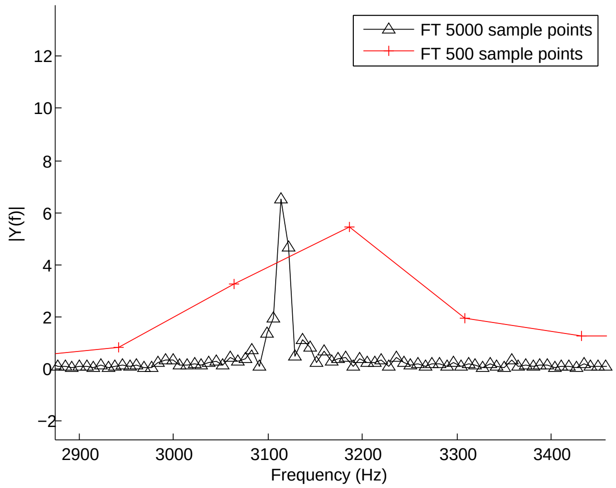

Figure 6. Zoomed in Single-Sided Amplitude Spectrum, FT stands for Fourier Transform

work with an array of sending emitters that relay the carrier waves at the points of maximum decay, so that the localization space can be extended by several orders of magnitude. This is a subject for future research.

4 Stochastic Radio Interferometric Positioning

This section describes our novel SRIPS positioning algorithm. It first analyzes and compares the optimization functions of RIPS and SRIPS. Then it describes the implementation of the localization algorithm of SRIPS. In the last subsection, it provides a typical example of the performance of RIPS and SRIPS.

4.1 RIPS versus SRIPS

RIPS ([8], [12] and [14]) estimates the position by sequentially minimizing two cost functions. First it calculates the q-ranges by minimizing the cost function described in Equation 8. This equation calculates the q-range that minimizes the squared difference between the phase offsets measured at N frequencies associated with one sender and receiver pair. This minimum is equal to the global minimum of the q-range error distribution. Then RIPS calculates the position by minimizing the cost function described in Equation 9. This latter cost function calculates the position by minimizing the squared difference between the calculated and estimated q-ranges. RIPS ([8], [12] and [14]) uses a Least Squares Method to calculate the position that best fits the measurements. In Radio Interferometric Positioning, the measured values are the phase offset measurements. However, the position calculated by Equations 8 and 9, does not minimize the squared difference between all measured and estimated phase offsets. The best fit to all measurements is obtained by:

$$\{\hat{x},\hat{y}\} = \arg\min_{\hat{x},\hat{y}} \sum_{j=1}^{M}\sum_{i=1}^{N}(\Delta\varphi_{i,j} - \widehat{\Delta\varphi_{i,j}})^2 \qquad (10)$$

Here, $\widehat{\Delta\varphi_{i,j}}$ is the estimated phase offset calculated by the position estimate. In other words, Equation 10 calculates the best position (fit) given the set of measurements and underlying propagation and statistical models, rather than first taking the minimum of a subset and substituting it in the total. Note that Equation 10 corresponds to the Maximum Likelihood Estimator when the phase offset measurements follow a normal distribution. This results in the following practical differences between RIPS and SRIPS:

- The original RIPS algorithm cannot cope with false global optimums and q-ranges located in local optimums in the q-range error distribution (e.g. [9]). This is because Equation 8 calculates each q-range separately on the basis of the phase offset measurements associated with that q-range. Equation 10 can cope with false global optimums and local optimums when the other phase offset measurements (those not associated with that specific q-range) can discriminate between the true and false optimums.
- The original RIPS algorithm cannot cope with varying q-range estimate precisions. It implicitly assumes that the precisions are equal for the estimated q-ranges, because it minimizes the equally weighted squared difference between the calculated and fitted q-ranges (Equation 9). SRIPS discriminates against such less precise or wider q-range error distributions, because these provide a constant contribution to the squared errors of Equation 10.

The advantage of Equations 8 and 9 over Equation 10 is that Equations 8 and 9 require significantly fewer computations using the analytic solver ([10]), while providing the required localization performance on the CC1000 (e.g. [12]).

4.2 SRIPS implementation

This section describes the implementation of the localization algorithm of SRIPS. Equation 10 is rewritten as:

$$\{\hat{x},\hat{y}\} = \arg\min_{\hat{x},\hat{y}} \sum_{j=1}^{M}\sum_{i=1}^{N}(\Delta\varphi_{i,j} - \widehat{\Delta\varphi_{i,j}})^2 \qquad (11)$$

$$= \arg\min_{\hat{x},\hat{y}} \sum_{j=1}^{M}\sum_{i=1}^{N}(d_{i,j} - \widehat{\text{q-range}}_j)^2 \qquad (12)$$

$$= \arg\min_{\hat{x},\hat{y}} \sum_{j=1}^{M} \text{ERROR}(\widehat{\text{q-range}}_j) \qquad (13)$$

SRIPS minimizes Equation 13. However, Equation 13 has a rapidly oscillating behavior as can be seen from the t-range error distribution shown in Figure 3 ($\sum_{j=1}^{M}\text{ERROR}(\widehat{\text{q-range}}_j)$). Therefore, the calculated errors for a given position estimate become unpredictable. SRIPS solves this problem by taking the envelope of the minimum of the q-range error distribution:

$$\text{ERROR}_{SRIPS}(\widehat{\text{q-range}}_j) = \min\left(\text{ERROR}([\widehat{\text{q-range}}_j - W, \widehat{\text{q-range}}_j + W])\right) \qquad (14)$$

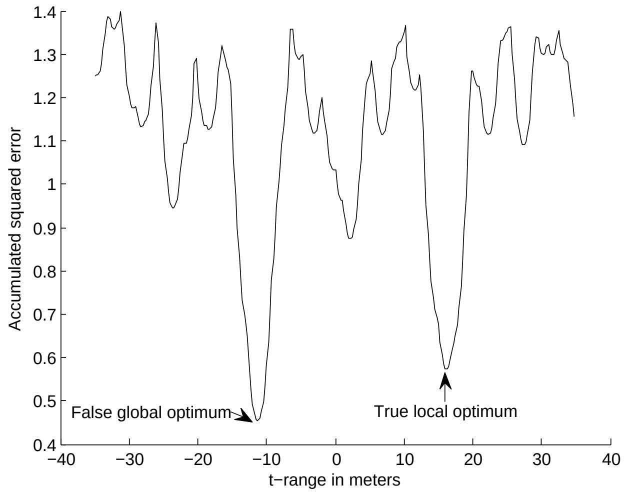

Figure 7. Smoothed t-range error distribution

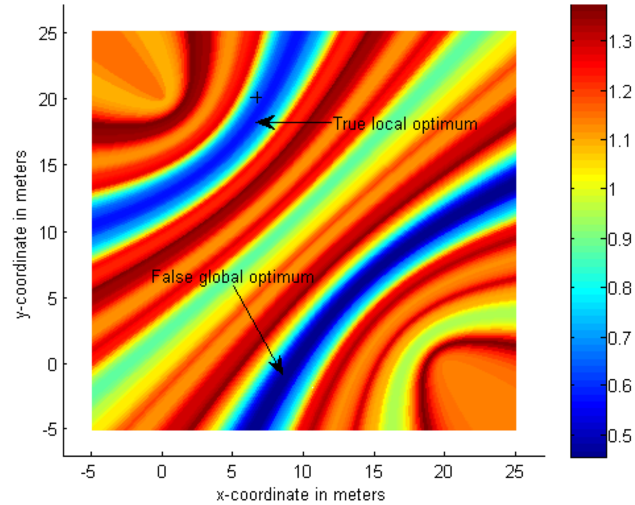

Figure 8. Smoothed t-range error distribution over the localization surface

Here W is a constant dependent on the frequency band used of the radio, in our case 2.4 GHz; $[\widehat{\text{q-range}}_j - W, \widehat{\text{q-range}}_j + W]$ is an interval; $\min\left(\text{ERROR}([\widehat{\text{q-range}}_j - W, \widehat{\text{q-range}}_j + W])\right)$ is the minimum value of Equation 7 over the specified interval. We numerically determine the value of W as 12.5 centimeters. We rewrite Equation 13 with the aid of Equation 14 as:

$$\{\hat{x}, \hat{y}\} = \arg\min_{\hat{x},\hat{y}} \sum_{j=1}^{M} \text{ERROR}_{SRIPS}(\widehat{\text{q-range}}_j) \qquad (15)$$

Figure 7 shows an example of the smoothed t-range error distribution using Equation 14. Figure 8 shows this smoothed t-range error distribution over the localization surface using Equation 15. In Figure 8, the cross is the true location of the target node; red represents large errors and blue represents small errors. Note that the two minimums shown in Figure 7 are represented by the two darker blue lines in Figure 8.

The SRIPS localization algorithm uses a grid-based Monte Carlo approach to minimize Equation 15. It accumulates the M smoothed q-range error distributions over the localization area represented by the Monte Carlo samples. In other words, it treats the q-ranges as stochastic variables to minimize Equation 15. We implemented SRIPS in Matlab, which post-processed all the data obtained by the CC2430 test bed. The main purpose of this implementation is to provide a tool for analysis; therefore, it does not minimize the computational or memory costs.

4.3 Typical Example of RIPS versus SRIPS

For reasons discussed in Section 5.5, our CC2430 platform test bed produces roughly 25% outliers in estimating the q-range, even at relatively long measurement times. Figure 9 shows an example where a t-range with a clear distinct minimum in its smoothed error distribution (blue curve) is combined with an outlier that has its global minimum about 30 meters away. These two t-range measurements belong to the same sender pair but have different receiver pairs. In other words, they share the same t-range (see Section 2.3). RIPS calculates the t-range based on these two measurements as the average of the global minimums of the two t-ranges, which is about $\frac{30}{2} \approx 15$ meters away from the true t-range. SRIPS averages the two error distributions as in Equation 15. This average is represented by the blue curve in Figure 10 and still has its global minimum at the true location.

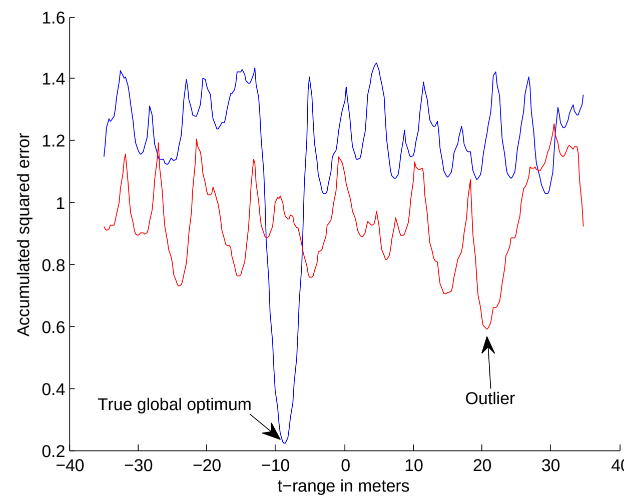

Figure 9. Two separate smoothed t-range error distributions (blue curve and red curve)

Figure 10 shows an example in one dimension (error as a function of the t-range). Note that SRIPS discriminates between true and false minimums in a similar way when all t-range error distributions are summed over the localization surface as defined in Equation 15. Figure 11 shows that SRIPS discriminates between the two minimums shown in Figures 7 and 8 by including the t-range error distributions from the other sender and receiver pairs. In this particular case, the resulting positioning error is 3 centimeters.

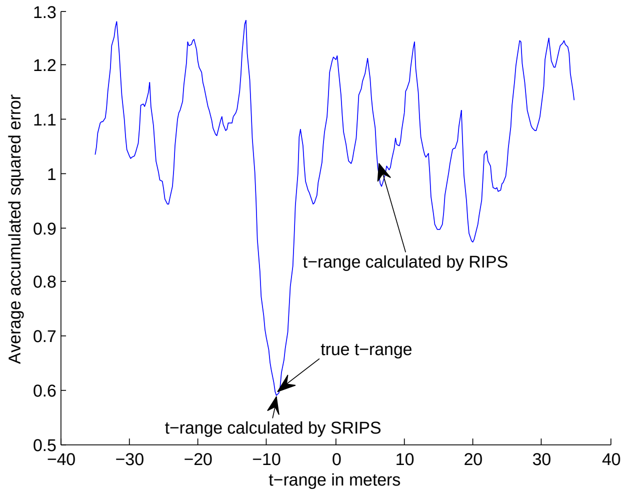

Figure 10. Average accumulated smoothed t-range error distribution as calculated by SRIPS

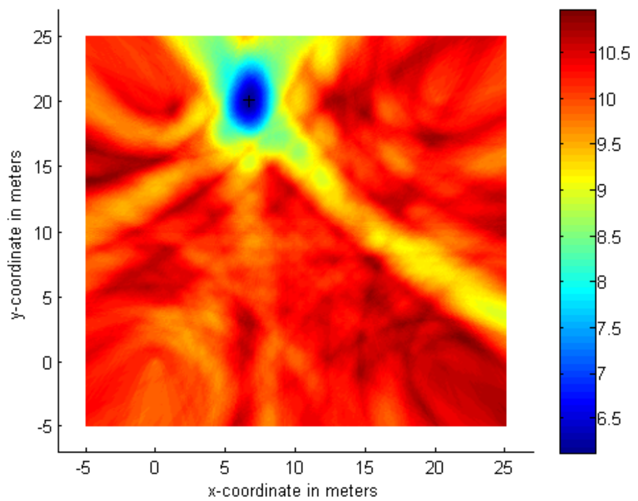

Figure 11. 2 s. total measurement time, error distribution over localization surface

Figure 12 illustrates the sensitivity of SRIPS to the total measurement time as defined by Equation 2. When reducing the total measurement time by a factor of 20, the error plot of Figure 12 becomes less discriminating. The positioning error increases to 60 centimeters. The advantage of these error distribution plots is that, at a glance, they give a quick impression of the accuracies of the position estimates.

5 Performance Evaluation

Section 5.1 describes the measurement set-up. Section 5.2 presents the performance evaluations of the RIPS and SRIPS positioning algorithms on the CC2430 platform with the various measurement settings of interest. Section 5.3 briefly discusses the amplitude filtering as applied by RIPS on the CC1000 platform. Section 5.4 compares the SRIPS results with ranged-based RSS results in a similar set-up. Section 5.5 summarizes the differences between SRIPS and RIPS on the CC2430 platform, and Section 5.6 compares the

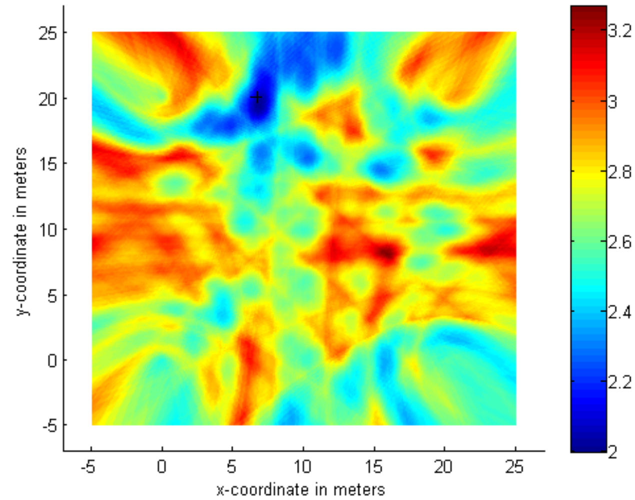

Figure 12. 100 ms. total measurement time, error distribution over localization surface

Figure 13. Measurement environment

RIPS and SRIPS results on the different platforms. Finally, Section 5.7 discusses the performance of SRIPS in indoor environments.

5.1 Measurement Set-up

The measurements were conducted in a 20×20 m^2 outdoor environment, shown in Figure 13, with six CC2430 radios (Figure 14). We used four CC2430 radios as reference nodes, which were located at the corners of the localization area; these reference nodes were fixed during and between measurement rounds. We used two CC2430 radios as target nodes; these target nodes measured the frequency beat signals at 12 different locations in a 4×4 grid. We determined the locations of the reference and target nodes with a measuring tape. We estimate the accuracy of the node placements as 10 cm. The receiving nodes measured 500 sample points per phase offset measurement in 8 ms per frequency over a total of 38 frequencies in a range of 2406 to 2480 MHz. All of the radios were placed at the same height of 1.5 m to minimize reflection noise (e.g. [9]). All individual sample points were sent to a computer and logged for post-processing. These measurements were performed over a period of two days with changing weather conditions. We

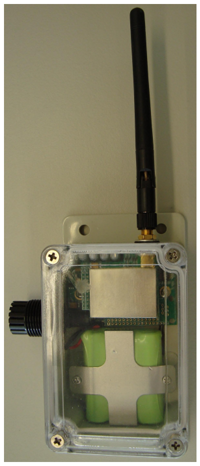

Figure 14. CC2430 radio

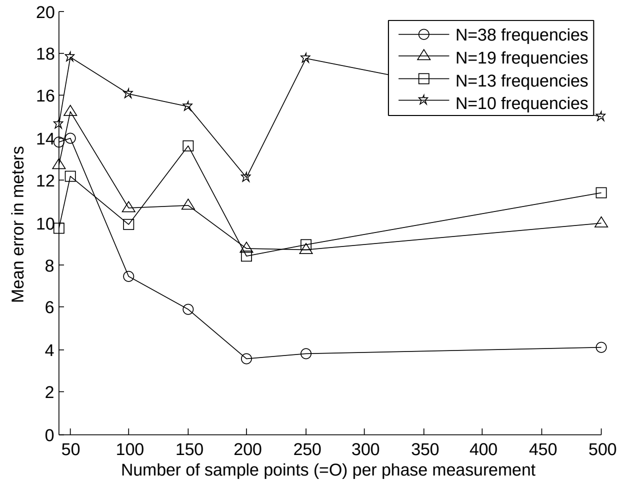

Figure 15. RIPS performance as a function of sample points (O) per phase measurement

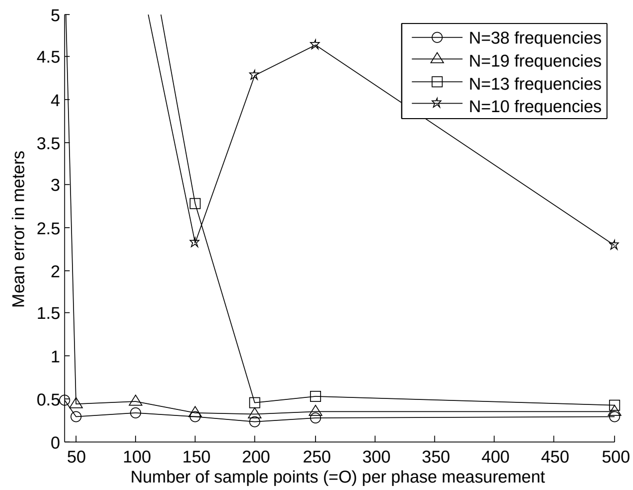

Figure 16. SRIPS performance as a function of sample points (O) per phase measurement

used this set-up because it is similar to the set-up described in [8].

This measurement set-up also performed 500 RSS measurements at the same 38 frequencies using individual sender rather than paired senders to compare the performance of RIPS and SRIPS with a typical RSS-based localization algorithm.

5.2 Performance Evaluation

This section analyzes the performance of SRIPS as a function of the total number of sample points to position one target node. We calculate this number and the measurement time using Equations 1 and 2. Figures 15 and 16 show the performance of the RIPS and SRIPS positioning algorithms as a function of the number of sample points per phase offset measurement. The horizontal axis represents the number of measured sample points per phase offset measurement. The four curves show the performance of the RIPS positioning algorithm using a different number of frequency measurements ($N = 38/19/13/10$) per sender pair ($M = 6$) and using the full frequency bandwidth of the available frequency band (2.406...2.480 GHz). This means that we increase the frequency hop length to $\{2,4,6,8\}$ MHz instead of decreasing the frequency bandwidth. We use this strategy because experiments show that it provides better results than the frequency-bandwidth strategy.

Figures 15 and 16 show that SRIPS converges faster and more accurately to the true locations than RIPS. These empirical numerical results indicate that beat signals on the CC2430 platform can be roughly collected over 0.8 ms (50 sample points) at 19 different frequencies and 6 sender pairs to have SRIPS yield reliable results on the order of 50 cm accuracy.

Figure 17 shows the number of measurable phase offset measurements as a function of the number of measured sample points per phase offset measurement. Not every measurement provides a measurable phase offset measurement due to noise or due to the fact that not every frequency beat is measurable given the measurement time and sampling rate (see Section 3.1). Figure 17 shows that the number of measurable phase offsets decreases when the number of sample points per phase offset measurement decreases below 200, especially from 100 to 50 sample points. Two sender pairs do not generate measurable frequency beats given the measurement time per phase offset measurement (50 samples, 0.8 ms), because one period of the generated beat signals is larger than the measurement time per phase offset measurement. So the effective total measurement time reduces from 0.09 to 0.06 s when sampling the beat signal with 50 sample points. This then implies that our nodes and environment must be static during the total measurement time of all beat signals, which is on the order of 0.1 seconds.

5.3 Amplitude Filtering

In this section we analyze whether the amplitude filter implemented on the CC1000 RIPS platform provides similar results on the CC2430 platform. The amplitude filter filters frequency beat signals with an amplitude smaller than a certain threshold ([8]). Figure 18 shows the performance of the

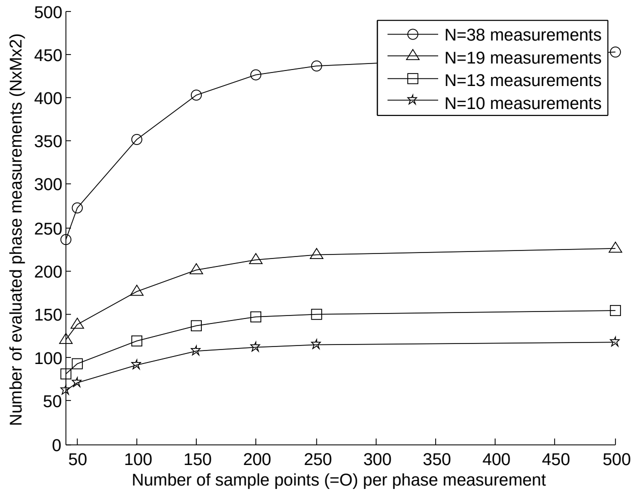

Figure 17. Number of valid phase measurements

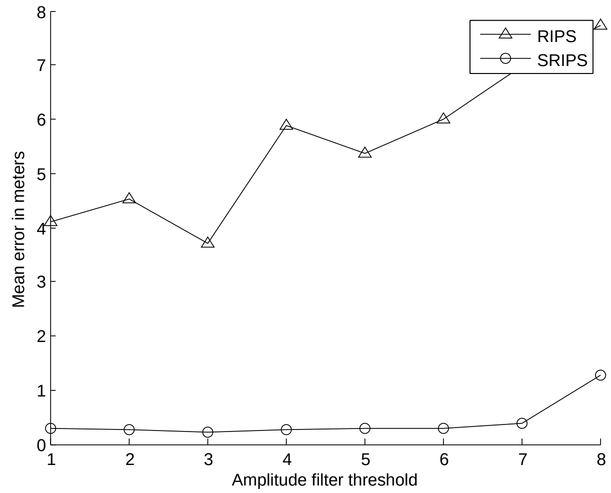

Figure 18. Performance as a function of the amplitude filter threshold

RIPS and SRIPS positioning algorithms as a function of this threshold. It shows that the amplitude filter does not significantly affect the performance of both algorithms on our platform, and that it can decrease the performance when it is set too high. This decreasing performance with increasing threshold filtering appears to be logical because increasing the threshold decreases the amount of evaluated phase measurements, which in turn decreases the localization performance.

5.4 Performance of Typical RSS-based Localization Technique

Recent studies ([18] and [19]) show that range-based localization ([6]) outperforms or provides similar results to other typical localization methods in line-of-sight environments (e.g. [5], [7]). Therefore, we evaluated the performance of the range-based localization algorithm described in [6]. First, we calibrated the propagation model ([4]) by the localization measurements. Then we used the same localization measurements for evaluating the performance. This ensures that [6] provides the best performance. [6] provides a mean error of ~ 6 meters, so SRIPS outperforms ranged-based RSS localization by an order of magnitude in this measurement set-up.

5.5 Discussion of CC2430 RIPS versus CC2430 SRIPS Results

Figures 15 and 16 show that:

- RIPS provides a localization accuracy of several meters rather than centimeters reported on the CC1000 platform.
- SRIPS outperforms the RIPS positioning algorithm in all cases investigated.

As described in Section 4.3, these differences are mainly caused by the fact that RIPS uses a least-squares method, which does not cope well with the 25% outliers produced in the q-range error distributions by our test bed.

As a first candidate for the cause of these outliers, we eliminated multipath effects, which is further explained in Section 5.7. Except for random spontaneous hardware effects, it is unlikely that the cause can be traced to the hardware. It is unlikely because the same hardware produces reliable measurements and outliers in the same locations at different times without any noticeable changes in the environment. As a possible cause for these outliers, we suspect the busy 2.4 GHz band, which carries a lot of traffic. All sorts of interferences, such as those from a nearby WiFi network, might have produced outliers in our test bed at unexpected times. The real world always has outliers and the good thing is that our localization algorithm dealt with them in an adequate way.

RIPS uses a least-squares method, which is known to be susceptible to outliers. If we use a Least-Absolute Deviation (LAD) rather than a least-squares optimization algorithm:

$$\{\hat{x},\hat{y}\} = \arg\min_{\hat{x},\hat{y}} \sum_{j=1}^{M} \left| q_{\text{est},j} - \text{q-}\widehat{\text{range}}_j \right| ,$$

it adequately manages the 25% outliers (~ 30 cm accuracy when evaluating all measurements). However, the performance of LAD decreases rapidly when decreasing the total number of sample points. LAD still provides decimeter accuracy when sampling 100 sample points per phase measurement at 38 frequencies. The error increases directly to several meters when we further decrease the total number of sample points. For example, when sampling 50 sample points per phase measurement at 38 frequencies, the localization accuracy decreases to ~ 9 meters while SRIPS still provides a localization accuracy of ~ 30 cm. Moreover, LAD needs a factor of six more measurements than SRIPS for providing decimeter accuracy, as we show in the next section.

References	Measurement time	Total sample points	Error	Area	Frequencies	Sender Pairs	Platform
[8]	88747 ms	NA	3 cm	324 m^2	13	240	CC1000
[9]	69518 ms	NA	10 cm	8000 m^2	13	188	CC1000
[11]	440 ms	43560	61 cm	7200 m^2	11	1	CC1000
[12]	1252 ms	33972	70 cm	780 m^2	22	2	CC1000
[14]	1024 ms	46080	70 cm	100 m^2	18	2	CC1000
SRIPS	60 ms	11400	50 cm	400 m^2	19	4	CC2430
RIPS	730 ms	136800	410 cm	400 m^2	38	6	CC2430
RIPS(LAD)	365 ms	68400	50 cm	400 m^2	38	6	CC2430
TYPICAL	2000 ms	76000	620 cm	400 m^2	38	4 senders	CC2430

Table 2. CC1000 RIPS and CC2430 SRIPS performance

5.6 CC1000 RIPS versus CC2430 SRIPS Results

Table 2 summarizes the results obtained by the CC1000 RIPS implementation and our CC2430 SRIPS implementation. This table consists of eight columns:

- **References** is the reference used or the name of the positioning algorithm evaluated.
- **Measurement time** is the total measurement time of one target node. We use Equation 2 to calculate the total measurement time.
- **Total sample points** is the total number of sample points used to position one target node. We use Equation 1 to calculate the total number of sampling points.
- **Error** is the mean of the positioning error.
- **Area** is the surface area of the localization surface.
- **Frequencies** is the number of frequencies used per sender pair.
- **Sender pairs** is the number of sender pairs used by RIPS and SRIPS.
- **Platform** is the radio platform used.

The first two rows show the performance of the RIPS network localization ([8] and [9]). This set-up characterizes itself by the relatively long sampling time (minutes) and high accuracy (centimeters). The remaining rows represent the performance using the infrastructure approach as described in Section 2.2. The difference between the two approaches clearly shows that the performance improves with increasing sampling time. [8]/[9] increase the performance by a factor of 10 compared to [11]/[12]/[14]/SRIPS at the cost of an increased sampling time by a factor of 250.

[11] differs from [12]/[14]/SRIPS in that it implements the target-as-sender instead of the target-as-receiver approach. It is difficult to compare the results between these two implementations because [11] calculates 55 q-ranges on the basis of one measurement round. The disadvantage of this approach is that the measurement time increases linearly with the number of target nodes, which is not the case with the target-as-receiver approach. In addition, [11] estimates positions of mobile nodes, while [12]/[14]/SRIPS estimate the positions of static nodes. Note that [12] and [14] also estimate the position of mobile nodes, but we leave these results out of the table. For completeness, we added the results of the Least-Absolute Deviation method described in Section 5.5.

Table 2 shows that SRIPS provides comparable results as reported by [11]/[12]/[14], while reducing the sampling time by a factor of 15 compared with a target-as-receiver implementation. A factor of 7 is explained by the seven times faster sampling rate of the CC2430. In general, SRIPS requires a factor of 3 fewer measurements compared with the CC1000 implementation. This is because RIPS requires q-range measurements with a relatively high precision (the global minimum is required to be at the true q-range). This high precision is obtained by increasing the measurement time per sender pair ([12] and [14]). SRIPS does not have this requirement as we have seen. Therefore, SRIPS can reduce the required measurement time per sender pair significantly, while increasing the number of sender pairs. This finally results in requiring a factor of 3 fewer measurements, excluding the seven times higher sampling rate available with our CC2430 radio platform.

5.7 Radio Interferometric Positioning in Indoor Environments

Both RIPS and SRIPS use the same propagation model that models the electromagnetic energy as scalar plane waves. This scalar approach neglects the polarization effects that start to play a role when the dimensions of the obstacles that meet these propagating waves become of the order of the wavelengths of the carrier waves. In addition, the propagation model does not account for multiple reflections, such as those that are commonly modeled in the area of image reconstruction. Hence, both RIPS and SRIPS cannot reliably localize in such environments. Our preliminary set of indoor measurements in a hall with metal frames confirms this. The results were especially unreliable when the nodes were near these metal frames. But we did obtain decimeter accuracy in those cases where only a few nodes suffered from interfering reflections. Hence, when multiple reflections are not dominating the measurements, SRIPS holds the promise of the ability to discriminate between those reflections.

6 Conclusion

In this study, we experimentally verified that it is possible to perform Radio Interferometric Positioning on commonly available radio platforms, such as the CC2430. Such radio platforms do not comply with the frequency tuning requirements of the existing RIPS implementation on the CC1000 platform.

Experiments on the CC2430 platform showed that the RIPS positioning algorithm does not provide decimeter accuracy, because it cannot cope with the outliers generated by the CC2430 platform. This paper introduced a novel RIPS algorithm, which we call Stochastic Radio Interferometric Positioning (SRIPS). SRIPS was shown to cope with the varying accuracies of the distance estimates without any calibration and with relatively short measurement times compared with the existing CC1000 RIPS implementation. Experiments in a 20×20 m^2 set-up showed that our SRIPS CC2430 implementation reduces the measurement time to less than 0.1 seconds, while providing an accuracy similar to the existing RIPS implementation on the CC1000 platform.

Future work consists of evaluating SRIPS on more radio platforms in more diverse environments.

7 Acknowledgements

We would like to thank our shepherd Ákos Lédeczi for his comments and guidance, and the anonymous reviewers for their constructive criticism. We are very grateful to Cindy Kleinfeld for editing the paper.

This paper describes work in part undertaken in the context of the TSP project. The TSP project has joint funding from the European Regional Development Fund of the European Union and the Dutch Provinces of Gelderland and Overijssel.

8 References

[1] M.Sargent, E.Willis Lamb, R.L.Fork: Theory of a Zeeman Laser I. Physical Review, vol. 164, Issue 2, pp. 436-449, December 1967.

[2] D.Rife, R.Boorstyn: Single tone parameter estimation from discrete-time observations. IEEE Transactions on Information Theory, Volume 20, Issue 5, pp. 591598, September 1974.

[3] Yi Xie and Yi-zun Wu: Zeeman laser interferometer errors for high-precision measurements. Applied Optics, Vol. 31, Issue 7, pp. 881-884 (1992)

[4] Hashemi H.: The indoor radio propagation channel, Proc. IEEE, July 1993, pp. 943- 996.

[5] P. Bahl and V. N. Padmanabhan: RADAR: An In-Building RF-Based User Location and Tracking System, in Proceedings of the 19th IEEE International Conference on Computer Communications (INFOCOM), March 2000.

[6] N.Patwari: Location estimation in sensor networks. Thesis of Neal Patwari at University of Michigan, 2005.

[7] K.Yedavalli, B.Krishnamachari, S.Ravula, and B.Srinivasan: Ecolocation: A sequence based technique for RF-only localization in wireless sensor networks. In IEEE IPSN 2005, April 2005.

[8] M.Maróti, P.Völgyesi, S.Dóra, B.Kusý, A.Nádas, Á.Lédeczi, G.Balogh, K.Molnár: Radio interferometric geolocation. SenSys 2005: pp. 1-12.

[9] B.Kusý, Á.Lédeczi, M.Maróti, L.G.L.T.Meertens: Node density independent localization. IPSN 2006: 441-448.

[10] B.Kusý, J.Sallai: Analytical solution for radio-interferometric localization of mobile sensors. ISIS technical report, ISIS-06-710. May 2006.

[11] B.Kusý, G.Balogh, J.Sallai, Á.Lédeczi, M.Maróti: InTrack: High Precision Tracking of Mobile Sensor Nodes. EWSN 2007: 51-66.

[12] B.Kusý, J.Sallai, G.Balogh, Á.Lédeczi, V.Protopopescu, J.Tolliver, F.DeNap, M.Parang: Radio interferometric tracking of mobile wireless nodes. MobiSys 2007: 139-151.

[13] B.Kusý, Á.Lédeczi, X.D.Koutsoukos: Tracking mobile nodes using RF Doppler shifts. SenSys 2007: 29-42.

[14] H.Wu, H.Chang, C.You, H.Chu, P.Huang: Modeling and optimizing positional accuracy based on hyperbolic geometry for the adaptive radio interferometric positioning system. International Symposium on Location- and Context-Awareness (LOCA 2007), Oberpfaffenhofen, Germany, September 2007, pp 228-244.

[15] J.Tian, H.Chang, T.Lai, H.Chu: SpinTrack: Spinning Infrastructure Nodes for Precise Indoor Localization. In the poster session of UBICOMP, September 2008.

[16] H.Chang, J.Tian, T.Lai, H.Chu, P.Huang: Spinning Beacons for Precise Indoor Localization. ACM SENSYS 2008, November 2008.

[17] B.J.Dil, P.J.M.Havinga: A Feasibility Study of RIP Using 2.4 GHz 802.15.4 Radios. MELT 2010, November 2010.

[18] B.J.Dil, P.J.M.Havinga: RSS-Based Localization with Different Antenna Orientations. Australian Telecommunication Networks and Applications Conference (ATNAC), 2010.

[19] B.J.Dil, P.J.M.Havinga: Calibration and Performance of RSS-based Localization Methods. Internet Of Things (IOT), 2010.

[20] CC1000. http://focus.ti.com/lit/ds/symlink/cc1000.pdf, 2011.

[21] CC2430. http://focus.ti.com/lit/ds/symlink/cc2430.pdf, 2011.

[22] MRF24J40. http://ww1.microchip.com/downloads/en/DeviceDoc/DS-39776b.pdf, 2011

[23] Low Power 2.4 GHz Transceiver for ZigBee, IEEE 802.15.4, 6LoWPAN, RF4CE and ISM Applications. http://www.atmel.com/dyn/resources/prod_documents/doc5131.pdf, 2011

[24] Product Brief JN5148 Module. http://www.jennic.com/files/product_briefs/JN5148-MO-PB_1v1.1.pdf, 2011