Towards Precise Indoor RF Localization

Akos Ledeczi*
Peter Volgyesi*
Janos Sallai*
Branislav Kusy†
Xenofon Koutsoukos*
Miklos Maroti‡

Abstract

Precise indoor localization of wireless nodes remains a challenge today. While there are radio-frequency (RF) methods that offer significant advantages, the balance between accuracy, range, and cost is suboptimal for many applications. Radio interferometry has been shown to be effective outdoors, however, its applicability indoors has not been demonstrated mainly due to its sensitivity to multipath. This paper presents a roadmap outlining how the method can be enhanced to advance the state-of-the-art in indoor RF localization.

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, Experimentation, Theory

Keywords

Sensor Networks, Doppler effect, Localization

Acknowledgments

This work was supported in part by NSF grant CNS-0721604 and ARO MURI grant W911NF-06-1-0076.

1 Introduction

Location-awareness is an important requirement for many mobile wireless applications. While there are many localization systems for outdoor use, fine-grain indoor localization is still a key missing piece for a range of applications such as asset tracking in a warehouse or locating emergency personnel in a disaster area.

Radio-frequency (RF) methods offer significant advantages in terms of range, scalability, deployment and maintenance. Although there are numerous noteworthy results, there still exist significant theoretical and practical challenges especially for providing high-precision, cost-effective, and scalable solutions indoors. The accuracy of RF location systems such as RADAR [1] and Horus [18] as well as numerous other radio signal strength (RSS) methods is typically meter-scale [16].

A new class of location systems has emerged with low-cost identification technologies such as RFID tags. While passive RFID tags are self-identifying, they are not self-locating, so they require additional hardware for localization, such as mobile robots, cameras, laser-range scanners, and RFID readers. Active RFID systems use self-powered tags to identify and locate objects. SpotON uses an aggregation algorithm for three dimensional location sensing based on RSS [7]. LANDMARC deploys multiple fixed RFID readers and reference tags as infrastructure [13]. The accuracy achieved is meter-scale and requires high density of the tags.

Ultra-wideband (UWB) systems are resistant to multipath propagation and have very good time domain resolution for localization and tracking [5]. UWB range measurements have demonstrated RMS ranging errors of 0.12 to 1.5 *m* [2]. A fine-grained localization system with an accuracy of about 20 *cm* was developed by Ubisense [15]. The disadvantage of UWB is that it requires high sampling rates ($> 1GHz$) and precise inter-node time synchronization ($< 1ns$) translating into higher cost. Also, the Federal Communications Commission (FCC) has mandated that UWB radio transmissions operate at limited transmit power restricting the maximum range. Effectively, most of the approaches consider maximum range of $10-20$ *m* [17, 4].

In contrast, radio interferometric localization has been shown to localize wireless sensor nodes with high accuracy even at long ranges without extra hardware [12]. However, the approach was developed for outdoor deployments and currently cannot meet the localization requirements of indoor applications without further theoretical and experimental research.

This paper presents a roadmap outlining how the radio interferometry method can be enhanced to advance the state-of-the-art in indoor RF localization. We discuss the main challenges that need to be addressed focusing on multipath propagation. We propose a group of promising methods that include the utilization of low carrier frequencies, redundancy in the frequency measurements and the number of anchors,

*Institute for Software Integrated Systems, Vanderbilt University, 2015 Terrace Place, Nashville, TN 37203, USA
†Stanford University, Palo Alto, CA, USA
‡Bolyai Institute, University of Szeged, Szeged, Hungary

HotEmNets'08, June 2–3, 2008, Charlottesville, VA, USA

incorporating RSS methods, and exploiting asymmetric architectures. Finally, we present preliminary results using a Software Defined Radio (SDR) platform. We show that the phase difference between two receivers can be measured without time synchronization simplifying the system design significantly.

2 Motivating application

One of the driving applications is precise tracking of assets in a warehouse. We envision a system that can locate boxes and pallets within a meter in 3D. While passive RFID systems are deployed today, the location information they provide is only the binary proximity function to a reader. Active RFID systems are gaining ground, especially for office building-like settings. For example, PanGo, is a commercial asset tracking system provides room-level granularity using 802.11 active RFID tags [14]. However, tag prices allow only the tracking of high-value assets, the access point infrastructure needs high density and the approach is not precise enough for a warehouse setting.

Here is how our system would work. The simple and inexpensive tag has a radio transmitter and a microcontroller. The tag is in low-power mode until a simple motion switch wakes it up. Once it stops moving, it waits for a random time, and then it broadcasts its ID and starts transmitting a set of sine waves in parallel for a short amount of time. The process may be repeated a few times and a simple carrier sense circuitry can be used to reduce collisions. Afterwards, the tag goes back to sleep. The warehouse is instrumented with a network of infrastructure nodes. As they receive messages from tags, they schedule ranging measurements. For each tag, a second transmitter is selected among the infrastructure nodes while all other nodes act as receivers. The entire ranging process takes a few *ms*, so up to one hundred tags could be processed every second. The localization algorithm determines the tag location and stores it in a database. Note that the usability of such technology will depend primarily on the achievable price, size, and battery life of the tags.

3 Radio interferometric positioning

Interferometry is a widely used technique in both radio and optical astronomy to determine the precise angular position of celestial bodies, as well as objects on the ground. The basic idea behind Radio Interferometric Positioning (RIPS) is to emit pure sine wave radio signals at two locations at slightly different frequencies [12]. The composite radio signal has a low beat frequency and its envelope signal can be measured as the received signal strength indicator (RSSI) signal. The relative phase offset of this signal between two receivers depends only on the distances between the two transmitters and two receivers and on the wavelength of the carrier signal (see Fig. 1 and Eq. 1). We call the ordered quadruple (A,B,C,D) of the distinct nodes a quad and the linear combination of distances $d_{AD} - d_{BD} + d_{BC} - d_{AC}$ for the quad the q-range d_{ABCD}. We denote the absolute phase offset of the RSSI signal measured by node X by φ_X, the relative phase offset between X and Y by $\varphi_{XY} = \varphi_X - \varphi_Y$, and the wavelength corresponding to the carrier frequency f of the radio signal by $\lambda = c/f$. Assuming that the frequency separation between the transmitters is relatively small and the

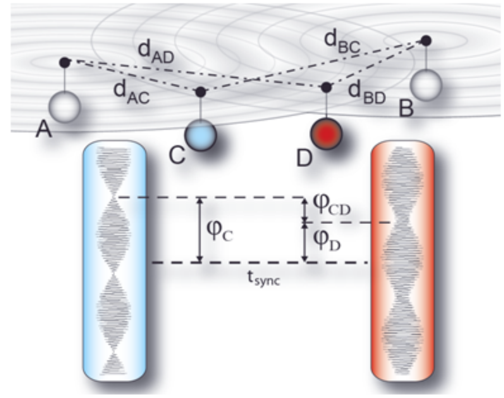

Figure 1. Radio Interferometric Ranging.

maximum pairwise distance between the four nodes is less than a few hundred meters, the following equation holds:

$$d_{ABCD} \; mod \, \lambda \; = \; \varphi_{CD} \, \frac{\lambda}{2\pi} \qquad (1)$$

The relative phase φ_{XY} can be measured by the receivers C and D. Note that a single φ_{XY} measurement does not yield a unique d_{ABCD} q-range because of the $(mod\,\lambda)$ term. However, we can measure φ_{CD} at different carrier frequencies and compute the q-range. If the network has at least 6-nodes, the q-ranges can be used to determine the relative node locations. RIPS has been implemented on Mica2 motes. One test was carried out in a rural area, a benign environment in terms of RF multipath propagation. 16 motes were deployed in an approximately $12000m2$ area with an average closest-neighbor distance of $35m$ and a maximum node distance of $170m$. The average localization precision was $4cm$, while the largest error was $12cm$, using three anchor nodes. Larger-scale experiments performed in urban environments with moderate multipath yielded less precision; depending on the elevation of the motes, their average distance, and the environment, the mean error was $10-50cm$.

Tracking systems typically locate mobile objects using a set of stationary infrastructure nodes at known locations. Using RIPS, successive phase measurements using several frequencies are taken at different geographic locations, hence the q-range d_{ABCD} is not constant during the ranging measurement. The inTrack system [8] tracks objects by defining the q-range to be an interval, rather than a single value, thus finding the location of the node in a region, rather than at a single point. This method can tolerate up to one half-wavelength change of the q-range during the ranging measurement. The mTrack system [11] compensates for the node mobility using the Doppler shift of the frequency of the interference signal. The fact that the Doppler shift of a $430MHz$ RF signal can be accurately measured on a Mica2 node illustrates the strength of radio interferometry. The frequency change corresponding to the movement of the node at $0.3m/s$ is about $0.4Hz$, a 10^{-9} ratio compared to the signal frequency. However, the generated interference signal is only a few hundred Hz, increasing this ratio to 10^{-4} and making the technique feasible on motes. In fact, the latest tracking system relies only on Doppler-shift measurements [10] eliminating the need for phase measurements and multiple carrier frequencies. However, this system cannot locate stationery nodes since they do not induce any Doppler shift.

4 Challenges

Multipath propagation. Multipath propagation introduces phase shifts in the measured signal. As it is exactly the phase that RIPS measures, this is a significant problem. Fig. 2 shows a simulated ideal and multipath distorted RIPS phase-field with up to 4 reflections in a $4 \times 5 \times 3m$ room at $430MHz$. The simulation accounts for reflections, the phase shifts, signal polarization, and the angle of incidence. The crosses indicate the transmitters, while the shades of gray show the absolute phase (white: $\varphi = 0$, black: $\varphi = 2\pi$) of the envelope signal. In the vicinity of the line between the two transmitters, the phase error is acceptable, while outside this area, it is quite large.

Measurement Efficiency. The process of a single RIPS measurement is relatively slow, and requires complex and precise coordination. A RIPS measurement using a single frequency takes around $0.04sec$ on a mote, which limits the maximum speed of mobile nodes. However, these limitations are due to the constraints of the hardware platform and not necessarily inherent to RIPS. For example, most of the complexity stems from the required "tuning" of the carrier frequencies of the two transmitters as the difference between the nominal and actual frequencies of the crystal oscillators can be larger than the frequency separation mandated by the maximum sampling rate on a mote. When not constrained to motes, more sophisticated methods become feasible.

Localization Algorithms. The search-based localization algorithm presented in [9] and [8], is centralized and computationally expensive. The tracking algorithm presented in [11] is much faster, but still requires a PC class computer. Localization using q-ranges has not been studied beyond our early work because of the unique RIPS feature of using four nodes to make a single ranging measurement. However, there is a vast literature on localization using pair-wise distances and the challenge is to extend those approaches to q-ranges.

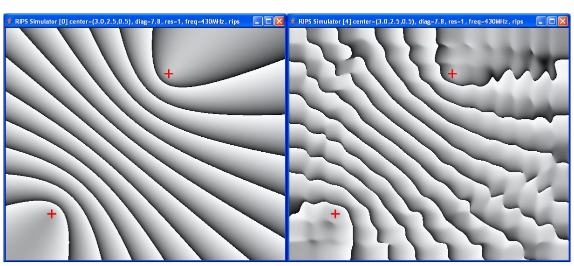

Figure 2. Ideal and multipath distorted phase fields.

5 Indoor localization

Low Carrier Frequency. Although it is impossible to avoid multipath effects indoors, their amount can be decreased. The main source of multipath is reflection. Reflections are caused by objects that have very large dimensions compared to the wavelength of the propagating wave. If the frequency of the carrier signal is decreased, so that the wavelength is larger than the typical walls and objects inside a building, reflections are going to diminish. For example, a $3MHz$ radio wave has $100m$ wavelength, so the signal will not reflect in typical buildings. In this case, diffraction and scattering become the main problems. Furthermore, as the RF signal passes through walls and objects, it incurs additional phase shifts depending on the material. However, if two RIPS receivers are located relatively close to each other then the two signals emitted by a transmitter will likely pass through mostly the same obstructions cancelling out these extra phase shifts.

There is another advantage of longer wavelength. If it is larger than twice the covered area then the $(mod\ \lambda)$ term disappears from Eq. 1, since all nodes are going to be located within the first period. Hence, it is enough to make the ranging measurements at a single frequency.

On the other hand, decreasing the carrier frequency introduces new problems. First, longer wavelengths require longer antennas. Second, there is only a limited set of lower frequencies that are not licensed already. Third, the longer wavelength may introduce near field effects close to transmitters. Finally, as indoor radio propagation is a very complex phenomenon, many detailed measurements need to be carried out in different environments before one can safely say that the approach is indeed feasible or not.

Redundant Carrier Frequencies. Some amount of multipath will be present regardless of the carrier frequency and the environment. One way that has already proved effective in moderate multipath environments is to make measurements at several frequencies. Fig. 3 shows a RIPS phase measurement for the same four nodes at 120 different frequencies The observed errors are as high as $3rad$. The environment was an outdoor urban area, but the nodes were deployed directly on the ground. Ground reflections significantly attenuated the direct line of sight signals, so the additional reflections from nearby buildings were causing significant phase errors. Interestingly enough, elevating the nodes by $1.5m$ decreased the errors by an order of magnitude because the effect of ground reflections was diminished.

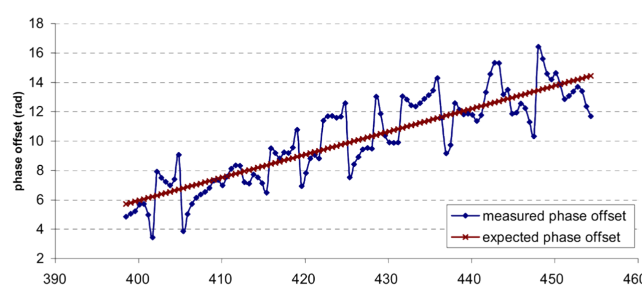

Figure 3. Phase measurements in a multipath environment.

Nevertheless, even the seemingly bad data in Fig. 3 is not hopeless. Using many data points, a straight line fitted on the measurements may have a small error. Of course, each measurement takes time, so there is a limit on the number of frequencies one can use.

Redundant anchor nodes. Having many anchor nodes enable us to include the same node in many quad configurations. Not only can then the error be decreased, but the erroneous q-ranges may be eliminated by exploiting an inequality between linear combinations of q-ranges, similarly to the triangle inequality for pair-wise ranges. When three out of the four nodes are anchors, we expect such an inequality to exist, since then the problem reduces to pair-wise distance differences. In general, this condition could be used to detect invalid q-ranges. As the number of possible q-ranges is significantly larger than that of linearly independent ones, anchor redundancy could provide a strategy to solve the localization problem even in the face of multipath.

Combining RIPS with RSS methods. If none of these approaches prove good enough to handle the extreme multipath environment of the inside of buildings, one can always turn back to proven methods and adapt them here. The fundamental idea of received signal strength (RSS) based localization systems is that the measured values are functions of the location of the participating nodes, and thus can be used to identify the location of tracked objects. This function can be deterministic [1], based on some path loss and fading model, or considered probabilistic utilizing a radio map constructed during an offline training phase [18]. RIPS is also applicable in the probabilistic case.

Existing RSS-based probabilistic localization algorithms can be augmented with RIPS data to achieve a more robust and precise localization in an indoor environment. The constructed radio map will contain the distribution of expected RSS values and RIPS phase correction information for all available channels. This offers the following advantages over traditional RSS methods: 1) with n infrastructure nodes, O(n2) essentially different RIPS phase measurements can be made as opposed to only O(n) RSS measurements, 2) the diversity of RIPS measurements over multiple channels is greater than that of RSS measurements, and 3) measurement error induced ranging errors (not localization errors) do not increase with the distance as was shown in [12]. We emphasize that this radio map based RIPS localization approach bypasses q-range calculations altogether and should work with significantly fewer radio channels.

Asymmetric Architecture. Unlike sensor networks where all nodes are equally resource constrained, typical indoor applications call for a different system architecture. While the tags to-be-tracked are resource constrained, the deployment area can be instrumented with a network of infrastructure nodes that are more powerful, e.g., with larger antennas, more memory and CPU cycles and wired power. This will enable them to carry out the more sophisticated processing techniques needed to increase the speed of ranging.

We can either increase the envelope frequency to shorten the time of an individual phase measurement or we can transmit multiple pairs of sinusoids at the same time (or both). Although the first approach seems to be a straightforward technological improvement, it might involve entirely new signal processing techniques. The most attractive feature of the second approach is that it enables simple and inexpensive transmitters with no time synchronization, therefore, it nicely supports asymmetric architectures.

Increased envelope frequencies require higher sampling rates and faster signal processing on the receivers. On the other hand, transmitter tuning via receiver feedback can be avoided if the frequency of the envelope signal is higher than the skew between the transmitters due to the frequency tolerance of the crystals ($30-50ppm$). Several improvements can be made beyond just increasing the sampling rate and processor speed. Signal processing in the frequency domain needs fewer captured periods, thus shorter measurement time. The required number of samples can be further decreased if oversampling and digital filtering is used to improve the signal quality. Also, the envelope signal can be extracted either by computing and filtering the squared signal (RSSI) or measuring the phases of the components directly at the intermediate frequency (IF).

Using multiple pairs of sinusoids simultaneously can drastically improve the measurement speed. The only requirement for the transmitters is the capability to generate and transmit multiple sinusoids in parallel, e.g., with $1MHz$ frequency separation. The receivers, on the other hand, will need more powerful signal processing capabilities. First, the radio frequency signal needs to be down-mixed to an IF signal, which is still several MHz wide. From this signal, we can extract the phases of each component in parallel using a high resolution FFT, or decompose the signal to individual pairs (by bandpass filters) and estimate the phases separately. Note that frequency tuning can be eliminated completely by setting the frequency of the second transmitter, a selected infrastructure node, in such a manner that even in the worst case of $50ppm$ error typical of crystal oscillators, the interference frequency will fall in the desired range.

6 Preliminary results

These ideas need to be experimentally validated as indoor radio propagation is notoriously hard to model. We selected Software Defined Radios (SDR) as our test platform since they can support a wide range of frequencies, they have powerful signal processing capabilities, and they are very flexible in terms of the supported radio protocols. They fit the asymmetric architecture nicely since they can play the role of the powerful infrastructure nodes or emulate the behavior of a simple, custom-designed tag.

We have implemented baseline RIPS ranging on the GNURadio platform [6] with USRP devices from Ettus Research [3]. We encountered a technical challenge that has implications for the final system. The two receivers in RIPS need to be time synchronized in order to compare the measured signal phases. However, we do not want to use the SDR for communication, only for RIPS ranging, since one would have to implement different protocols depending on the frequency band. Also, the wireless medium would need to be shared between localization and communication requiring a coordination protocol. We could not use standard protocols (e.g., IP over WiFi) for time synchronization between the SDR nodes because these methods do not provide the required accuracy, and the synchronization between the nodes and the radio peripherals is not solved.

Hence, our goal became to measure the phase difference between the two receivers without time synchronization simplifying the system design. The underlying idea is to have one of the transmitters embed an indicator in its signal marking a common point in time for both receivers. Then the receivers can measure the phase of the interference signal relative to this common reference.

Sending a simple impulse signal is not optimal as the limited bandwidth of the baseband digital channel restricts the emission of rising edges with arbitrarily steep slope. The emitted signal is a Hamming windowed linear frequency modulated (chirp) signal that is commonly used in radar applications. When the received signal goes through an appropriate matched filter, the output will be a sharp pulse. The matched filter is realized as an FIR filter on the SDR. The frequency of the chirp signal sweeps through a $30kHz$ band, $100kHz$ under the lower of the carrier frequencies around $450MHz$. Note that a bandpass filter is also applied to the already down-converted signal to filter out the carrier frequencies and noise before the matched filter is applied.

The procedure for RIPS ranging is performed as follows. One of the transmitters is always on, as the Received Signal in Fig. 4 shows initially. The second one emits the chirp signal. The TimesyncSignal shows the output of the matched filter at the receivers. An adaptive threshold detector observes the pulse and registers the reference point. After a $100\mu s$ pause, the transmitter starts emitting the carrier signal at a frequency slightly different than the first one. The signal strength of the Received Signal is the RIPS Signal shown in Fig. 4. After waiting for another $100\mu s$, the receivers start recording a 4096-sample long buffer and run an FFT to measure the phase of the signal. The receivers send their measured phases to a central processing computer that computes the relative phases. The procedure is then repeated for multiple carrier frequencies. The process can be optimized by having just one chirp signal for ranging measurements at multiple carrier frequencies. If the whole process takes $10ms$ and the interference frequency is as high as $1kHz$, a $50ppm$ clock error could only cause a maximum error of 0.2 degrees in the measured phase.

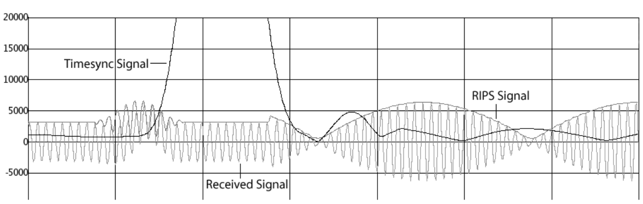

Figure 4. Measured RIPS signals on an SDR receiver.

It is hard to quantify the accuracy of this implicit time synchronization since it includes the accuracy of the phase measurement and the actual physical phenomena effecting the phases, such as multipath. Instead, we measured the jitter by repeating the procedure once a second 1000 times. The average jitter obtained was less than $1\mu s$, while the maximum was $2\mu s$, which at $1kHz$ frequency separation, yields less than 1 degree error.

7 References

[1] P. Bahl and V. N. Padmanabhan. Radar: An in-building RF-based user-location and tracking system. *Proc. IEEE INFOCOM*, 2000.

[2] N. Correal, S. Kyperountas, Q. Shi, and M. Welborn. An uwb relative location system. *In Proc of UWBST*, 2003.

[3] Ettus Research LLC, manufacturer of the Universal Software Radio Peripheral. http://www.ettus.com/.

[4] R. Fontana, E. Richley, and J. Barney. Commercialization of an ultra wideband precision asset location system. *In Proc. of UWBST*, 2003.

[5] S. Gezici, Z. Tian, G. B. Biannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios. *IEEE Signal Processing Magazine*, 2005.

[6] GNU Radio website. http://gnuradio.org/.

[7] J. Hightower, R. Want, and G. Borriello. SpotON: An indoor 3D location sensing technology based on RF signal strength. *UW CSE 00-02-02*, Feb. 2000.

[8] B. Kusý, G. Balogh, A. Lédeczi, J. Sallai, and M. Maróti. inTrack: High precision tracking of mobile sensor nodes. *In Proc. of EWSN*, 2007.

[9] B. Kusý, G. Balogh, P. Völgyesi, J. Sallai, A. Nádas, A. Lédeczi, M. Maróti, and L. Meertens. Node-density independent localization. In *In Proc. of IPSN/SPOTS*, 2006.

[10] B. Kusý, A. Ledeczi, and X. Koutsoukos. Tracking mobile nodes using rf doppler shifts. In *Proc. of SenSys*, 2007.

[11] B. Kusý, J. Sallai, G. Balogh, A. Lédeczi, V. Protopopescu, J. Tolliver, F. DeNap, and M. Parang. Radio interferometric tracking of mobile wireless nodes. *In Proc. of MobiSys*, 2007.

[12] M. Maróti, B. Kusý, G. Balogh, P. Völgyesi, A. Nádas, K. Molnár, S. Dóra, and A. Lédeczi. Radio interferometric geolocation. *In Proc. of ACM SenSys*, Nov. 2005.

[13] L. Ni, Y. Liu, Y. C. Lau, and A. Patil. Landmarc: indoor location sensing using active rfid. *In Proc. of PerCom*, 2003.

[14] PanGo. `http://www.pangonetworks.com`.

[15] UbiSense. `http://www.ubisense.net`.

[16] K. Yedavalli, B. Krishnamachari, S. Ravula, and B. Srinivasan. Ecolocation: a sequence based technique for rf localization in wireless sensor networks. *In Proc. of IPSN*, 2005.

[17] D. Young, C. Keller, D. Bliss, and K. Forsythe. Ultra-wideband (uwb) transmitter location using time difference of arrival (tdoa) techniques. *Signals, Systems and Computers, Asilomar Conference on*, 2003.

[18] M. Youssef and A. Agrawala. The horus wlan location determination system. In *In Proc. of MobiSys*, 2005.