
RF angle of arrival-based node localisation

Isaac Amundson*, Janos Sallai,
Xenofon Koutsoukos and Ákos Lédeczi

Institute for Software Integrated Systems (ISIS),
Vanderbilt University,
Nashville, TN, USA
Email: iamundson@gmail.com
Email: janos.sallai@vanderbilt.edu
Email: xenofon.koutsoukos@vanderbilt.edu
Email: akos.ledeczi@vanderbilt.edu
*Corresponding author

Miklós Maróti

Bolyai Institute,
University of Szeged,
Szeged, Hungary
Email: mmaroti@math.u-szeged.hu

Abstract: Several localisation algorithms exist for wireless sensor networks that use angle of arrival measurements to derive node position. However, there are limited options for actually obtaining the angle of arrival using resource-constrained devices. In this paper, we describe a technique for determining node bearings based on radio interferometric angle of arrival measurements from multiple anchor nodes to any number of target nodes at unknown positions. Least squares triangulation is then used to estimate node position. The position estimation is carried out by the node itself, hence the method is distributed, scalable, and fast. Furthermore, this technique does not require modification of the mote hardware because it relies only on the radio. Experimental results demonstrate that our approach can estimate node bearing with an average accuracy of 3° , and node position with sub-metre accuracy in approximately 1 s.

Keywords: sensor networks; localisation; radio interferometry; angle of arrival.

Reference to this paper should be made as follows: Amundson, I., Sallai, J., Koutsoukos, X., Lédeczi, Á. and Maróti, M. (2011) 'RF angle of arrival-based node localisation', *Int. J. Sensor Networks*, Vol. 9, Nos. 3/4, pp.209–224.

Biographical notes: Isaac Amundson received his PhD in Computer Science from Vanderbilt University, where he was a research assistant at the Institute for Software Integrated Systems. His research involves mobile wireless sensor network coordination, including time synchronisation, localisation, and robotic navigation.

Janos Sallai is a Research Scientist with Institute for Software Integrated Systems, Vanderbilt University. His research interests include execution environments and languages for low-power devices, applied sensor networks, as well as localisation and tracking in wireless sensor networks. He received his PhD in Computer Science from Vanderbilt University (2008) and his MSc in Computer Science from the Budapest University of Technology and Economics (2001).

Xenofon Koutsoukos is an Associate Professor in the Department of Electrical Engineering and Computer Science at Vanderbilt University. He received his PhD in Electrical Engineering from the University of Notre Dame. His research work is in the area of cyber-physical systems with emphasis on formal methods, distributed algorithms, diagnosis and fault management, and adaptive resource management. He received the NSF Career Award in 2004, the Best Paper Award in SenSys 2007, the Honourable mention for the Jean-Pierre Le Cadre Best Paper Award in Fusion 2011, and the Excellence in Teaching Award in 2009 from the Vanderbilt University School of Engineering.

Ákos Lédeczi is an Associate Professor in Vanderbilt University's Department of Electrical Engineering and Computer Science and a Senior Research Scientist at its Institute for Software Integrated Systems. His research interests include model-integrated system development and Wireless Sensor Networks (WSNs). His group is responsible for the first WSN-based

countersniper system, FTSP, the most widely used time synchronisation protocol in WSNs as well as radio interferometric localisation of wireless sensor nodes. He has a PhD in Electrical Engineering from Vanderbilt University.

Miklós Maróti obtained his PhD in Mathematics at Vanderbilt University in 2002. He has worked on the algorithmic complexity of various algebraic problems, the constraint satisfaction problem, modelling of embedded software, and routing, time synchronisation and localisation in wireless sensor networks. Currently he is an Associate Professor of Mathematics at the University of Szeged.

1 Introduction

Wireless sensor network (WSN) localisation has matured to the point where we can achieve accuracy in the sub-metre range. However, most techniques require sophisticated sensors, a heavy-weight infrastructure, extensive PC processing, and/or long localisation latencies. An approach that does not require additional hardware runs on the nodes themselves, and is fast enough to support mobility as well as stationary localisation, would be a step forward in this domain.

Localisation typically involves the transmission of a signal by at least one node, followed by signal processing on other participating nodes (Hightower and Borriello, 2001). Radio Frequency (RF), acoustic, infrared, and visible light are commonly used modalities in WSNs for determining spatial relationships. The choice of signal modality is important for accurate localisation, and depends on node hardware, the environment, and the application. WSNs are developed to provide inexpensive wide-area observation capability, and therefore it is generally undesirable to add additional hardware to the sensor board, because this increases cost, weight, and resource utilisation. Since all wireless sensor nodes have onboard radio hardware, RF propagation has become a popular signal modality for localisation. RF signal properties such as strength (Bahl and Padmanabhan, 2000), phase (Maróti et al., 2005), and frequency (Chang et al., 2008) have all been analysed to derive range or bearing for position estimation.

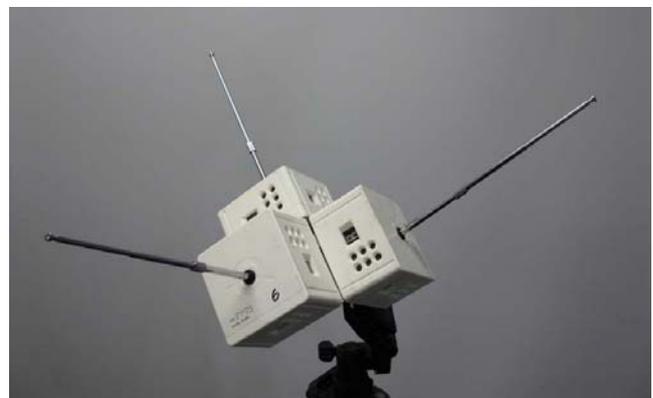
Radio interferometric positioning has emerged as an accurate method for sensor node localisation (Kusý et al., 2007b; Amundson et al., 2008; Chang et al., 2008; Lucarelli et al., 2008). The Radio Interferometric Positioning System (RIPS) demonstrated that the relative position of nodes can be determined accurately by only using their radios (Maróti et al., 2005). Although the resource-constrained nodes could not sample the pure RF signal fast enough, they could process the lower-frequency envelope of the beat signal that resulted from two high-frequency signals interfering. The difference in signal phase measured by two other nodes is a linear combination of the distances between the transmitters and receivers. Using radio interferometric ranging in this manner, relative node positions could be obtained in networks containing as little as six nodes.

There are several disadvantages associated with RIPS. To actually localise the nodes, phase differences must be obtained over several frequencies due to the modulo ambiguity caused by the phase wrapping to zero at 2π . To achieve precise results, RIPS relies on interferometric measurements using a large

number of combinations of four nodes. These measurements are carried out in a sequential manner. Finally, the data need to be routed back to the base station because the localisation algorithm runs on a resource-rich PC-class device. The entire process is relatively slow; for example, the results reported in (Maróti et al., 2005) took 50 min to obtain.

In this paper, we present TripLoc, a rapid distributed localisation method that uses radio interferometry, but overcomes many of the drawbacks of RIPS. Although we employ the same underlying radio interferometric measurement technique, our approach is otherwise original and based entirely on new work. The basic idea behind TripLoc is to group together three of the four nodes involved to form an antenna array and act as an anchor node. The two transmitters and one of the receivers are arranged in such a manner that their antennas are mutually orthogonal to eliminate parasitic antenna effects (see Figure 1). Since all the antennas are within half the wavelength, the modulo ambiguity is also eliminated. Each target node that needs to be localised acts as the second receiver in this scheme. The measured phase difference between the receiver of the array anchor node and the target node constrains the location of the latter to a hyperbola. This can be considered a bearing estimate, assuming that the node is not too close to the array. Hence, two such anchor arrays (or three non-collinear ones in degenerate cases) are enough to localise any number of nodes within range.

Figure 1 TripLoc antenna array implementation using three XSM motes. The name TripLoc comes from the fact that a triplet of nodes constitutes an anchor node array in our localisation scheme



This paper presents several new contributions to radio interferometric localisation.

- 1 We describe a radio interferometric angle-of-arrival-based technique for determining node bearing and position.
- 2 We provide a detailed analysis that shows our bearing estimation and localisation technique is robust to measurement noise and approximation error.
- 3 We design a real-world implementation using COTS sensor nodes, in which bearing estimation and triangulation is performed entirely on the resource-constrained nodes without additional hardware support.
- 4 We present experimental results that show our approach can rapidly and accurately estimate node bearing and position.

The remainder of this paper is organised as follows. In Section 2 we discuss other angle-of-arrival and radio interferometric localisation techniques for WSNs. Section 3 describes our system design, followed by an error analysis in Section 4. In Section 5, we describe our implementation on a real-world WSN platform. In Section 6, we evaluate our system based on experimental results. We conclude in Section 7.

2 Related work

The past decade has seen the emergence of numerous methods for sensor network localisation. These methods take different approaches, and use different propagation media, such as acoustic (Harter et al., 1999; Priyantha et al., 2000), radio frequency (Bahl and Padmanabhan, 2000; Maróti et al., 2005), and infrared (Want et al., 1992). The two general approaches to localisation are ranging, in which a node estimates its distance to multiple reference points; and bearing estimation, in which a node estimates its bearing to multiple reference points. Trilateration and triangulation-based techniques can then be respectively used to determine node position.

2.1 Angle of arrival

Node position is typically determined by estimating range or bearing to known positions. TripLoc employs the latter approach, whereby bearings from several anchor nodes are estimated, and then used to derive location. Although several techniques exist for determining node position based on bearing information (Esteves et al., 2003; Niculescu and Nath, 2003; Biswas et al., 2006; Rong and Sichitiu, 2006; Ash and Potter, 2007), there are few options for actually measuring signal AOA in WSNs. Currently available methods for bearing estimation require a heavy-weight infrastructure (Nasipuri and el Najjar, 2006), rotating hardware (Römer, 2003; Chang et al., 2008), directional antennas (Ash and Potter, 2004), and/or expensive and sophisticated sensors (Friedman et al., 2008).

Angle of arrival can be used in different ways for spatial coordination. Triangulation, for example, is the process of determining the position of an object from the bearings of known reference positions. Two such reference positions (or three non-collinear ones in degenerate cases) are enough to localise any number of nodes within range. Esteves et al.

(2003) gave a method to determine position based on the angular separation (the difference in bearings) between beacons. Other angle of arrival positioning approaches have been developed, including multi-angulation using subspace methods (Ash and Potter, 2007), anchor bearing propagation (Niculescu and Nath, 2003), and semi-definite programming (Biswas et al., 2006). Bearing estimates can also be useful when anchor positions are unknown. Altun and Koku (2005) and Bekris et al. (2004) presented mobile robot navigation methods for arriving at a target position by only observing angular separation between two pairs of landmarks.

A broad spectrum of acoustic beamforming techniques have been proposed to find the angle of incidence of a signal at an array of sensors. The most common techniques include delay-and-sum beamforming, Capon beamforming (Capon, 1969), MUSIC (Schmidt, 1986), ESPRIT (Roy et al., 1986) and min-norm (Kumaresan and Tufts, 1983) algorithms. Since the time of flight of the signal from the source to sensors in the array varies based on their pairwise distances, sensors receive the signal with different phases. While all of these methods compute the bearing of the source from the data streams sampled at the individual sensors, they differ greatly with respect to their angular resolution as well as their computational requirements.

The Cricket Compass is a device which uses ultrasound to determine orientation with respect to a number of ceiling-mounted beacons (Priyantha et al., 2001). Two receivers are mounted a few centimetres apart on a portable device, and the phase difference of the ultrasonic signal is measured to determine bearing. Although both the Cricket Compass and our approach measure signal phase difference to derive AOA, the two systems use different hardware, signal modalities, phase disambiguation techniques, and bearing derivation algorithms. The Cricket Compass has an accuracy of between 3° and 5° , depending on the orientation of the compass. The bearing accuracy of the Cricket Compass is comparable to TripLoc; however, Cricket requires customised hardware that includes an ultrasonic transceiver to compute bearing. Furthermore, ultrasound signal attenuation is much higher than radio, limiting the range of the Cricket Compass.

2.2 Radio interferometry

In Section 1, we described radio interferometric ranging. Here, we review other localisation techniques for WSNs that use radio interferometry, the RF method we employ in TripLoc. Although the techniques described below all employ the same measurement mechanism, the way the measurement data are used for localisation vary significantly and are distinct from the localisation method used in TripLoc.

mTrack is an extension of RIPS for tracking mobile nodes that makes several substantial improvements, such as incorporating the effects of mobility into the localisation algorithm (Kusý et al., 2007b). This is important because the position of the mobile node will change during the ranging measurement, and the RF signal will undergo Doppler shift due to the movement of the mobile node relative to the transmitters. In addition, mTrack computes the location of the tracked node more efficiently by

computing the intersection of two hyperbolic curves. Each curve represents an equation that constrains the location of the tracked node to some point on each hyperbola. By finding the intersection of two such curves, the position of the tracked node can be determined. In a 30×30 m field, mTrack reported an accuracy of 0.94 m, 0.18 m/s, and 11.61° in position, speed, and heading, respectively. Although TripLoc is unable to directly compute speed and heading, it performs localisations an order of magnitude faster, does not require centralised processing, and its position accuracy is slightly better than mTrack.

dTrack is a system for tracking mobile nodes that uses the same underlying principle of radio interferometry as RIPS, but measures signal frequency rather than phase (Kusý et al., 2007a). Tracking is accomplished by observing the Doppler shift in frequency that occurs when the source of a transmitted signal is moving relative to an observer. Because Doppler shift is determined by the relative velocity of the transmitter and receiver, absolute translational velocity of a mobile node can be determined using a priori knowledge of the transmitted frequency and the frequencies observed by stationary sensor nodes at known positions. Measurement noise was addressed by employing an Extended Kalman Filter (EKF). However, it was observed that the EKF did not respond well to sudden manoeuvres of the mobile node, and so it was combined with a constrained nonlinear least squares algorithm, which would update the EKF state if a manoeuvre was detected, thereby reducing the filter convergence time significantly. In a 50×30 m field, dTrack reported a combined CNLS/EKF location accuracy of between 1.4 m and 2.5 m in the best and worst case, respectively. The accuracy of the speed and bearing estimates were 0.1–0.4 m/s and $7\text{--}18^\circ$, respectively, and the location accuracy improvement of the CNLS/EKF algorithm over just the EKF was 9–36%. The main drawback with dTrack, compared with TripLoc, is that the sensor node must be moving in order to determine its position. Furthermore, the speed of the node must fall between fairly tight lower and upper bounds, determined by the carrier frequency and RSSI sampling rate of the mote radio. On average, dTrack position error is half a metre greater than TripLoc.

SpinLoc also uses the Doppler shift in the beat signal to determine position, however, unlike dTrack, the transmitter rotates about a fixed point at a constant speed (Chang et al., 2008). This has the advantage that the system will still work when nodes in the sensing region are stationary (i.e. translation is not required, as it is for dTrack). Spinning beacons and an assistant node at known positions transmit an interference signal with a predetermined beat frequency. A target node at an unknown position, as well as a reference infrastructure node, receives the signal. Because the beacon is spinning, the signal will undergo a Doppler shift, based on the instantaneous relative speed of the beacon and receiver. Using a Doppler angulation algorithm, the bearings of the target and reference node from the beacon are computed. Triangulation is then used to determine the position of the target node. What is perhaps most impressive about this system is that SpinLoc was able to obtain accurate results in a parking garage, a high multipath environment. In a 8×10 m area, the median position error for

SpinLoc was 39 cm, and in 90% of cases the error was 70 cm or less. SpinLoc performs slightly better than TripLoc; however, it requires specialised rotating hardware.

Lucarelli et al. (2008) follow a probabilistic approach to localisation, based on non-parametric belief propagation. The model was designed to take advantage of specific properties of radio interferometric ranging, such as the quad-range geometry of the nodes participating in the measurement. Although the approach is resource intensive, it was developed primarily to handle the uncertainty of wave propagation in multipath environments, as well as other causes of measurement noise. Using the same phase measurement data that were used by Maróti et al. (2005), the authors present simulation results of less than 15 cm average position error.

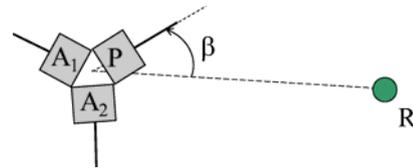
An adaptive radio interferometric positioning method was developed by Chang et al. (2007) that improves upon the positional accuracy of RIPS by only considering pairs of transmitters (for interference signal generation) that will give the best results based on the estimated position of the target node. An Estimation Error Model was developed which optimally selects the transmitters. Adaptive RIPS demonstrated an improvement over static RIPS of 47% and 61% for single and multiple targets, respectively.

3 System design

3.1 Radio interferometric measurements

Our system consists of multiple antenna array anchors at known positions and cooperating target sensor nodes at unknown positions. The arrays contain three nodes, a primary (P) and two assistants (A_1 , A_2), as shown in Figure 2. We assume that the position of the midpoint of the array is known, as well as the distance between the antennas in the array. At a predetermined time, the primary, P , and one of the assistants, A_1 , transmit a pure sinusoidal signal at slightly different frequencies, which interfere to create a low-frequency beat signal whose phase is measured by the other assistant in the array, A_2 , and a receiver node, R , at an unknown position. Such a measurement is termed a radio interferometric measurement (RIM).

Figure 2 Array containing a primary node (P) and two assistant nodes (A_1 , A_2). A target node (R) computes its bearing (β) from the array (see online version for colours)



The difference in phase, $\Delta\varphi = \varphi_R - \varphi_{A_2}$, measured by receiver nodes R and A_2 is a linear combination of the distances between the transmitters and receivers (Maróti et al., 2005),

$$\Delta\varphi = \frac{2\pi}{\lambda} (d_{P A_2} - d_{A_1 A_2} + d_{A_1 R} - d_{P R}) (\text{mod } 2\pi),$$

where λ is the wavelength of the carrier frequency, d_{PR} is the distance between the primary node and target receiver node, d_{A_1R} is the distance between the assistant transmitter node and the target receiver node, and d_{PA_1} , d_{PA_2} , and $d_{A_1A_2}$ are the respective distances between all pairs of nodes in the array. Note that the nodes in the array are equidistant from each other, and therefore $d_{PA_2} - d_{A_1A_2} = 0$, so the phase difference can be simplified:

$$\Delta\phi = \frac{2\pi}{\lambda}(d_{A_1R} - d_{PR})(\text{mod } 2\pi). \quad (1)$$

We denote the distance difference $d_{A_1R} - d_{PR}$ by d_{A_1PR} and refer to it as a *t-range*. From equation (1), we see that if $-\frac{\lambda}{2} < d_{A_1PR} < \frac{\lambda}{2}$, the phase difference will fall in the interval $(-\pi, \pi)$. When this is *not* the case, the possible range of $\Delta\phi$ will exceed 2π , which results in a modulo 2π phase ambiguity. To avoid this ambiguity, we would like the maximum possible distance difference to be less than $\frac{\lambda}{2}$. The maximum distance difference will occur when the receiver node is collinear with the transmitters P and A_1 . d_{A_1PR} then corresponds to the distance between the primary and assistant nodes. Therefore, to eliminate the modulo 2π phase ambiguity, we require the distance between antennas in the array to be less than half the wavelength of the carrier frequency.

Having removed the modulo operation, we can rearrange equation (1) so that known values are on the right-hand side:

$$d_{A_1PR} = \frac{\Delta\phi\lambda}{2\pi} \quad (2)$$

3.2 Bearing approximation

The *t-range* d_{A_1PR} defines an arm of a hyperbola that intersects the position of node R , and whose asymptote passes through the midpoint of the line $\overline{A_1P}$, connecting the primary and assistant nodes. Figure 3 illustrates such a hyperbola with foci A_1 and P . The absolute value of the distance differences between the foci and any point on a hyperbolic arm is constant, formally defined as

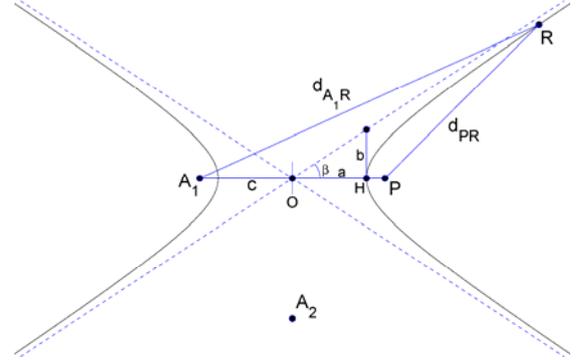
$$\frac{x^2}{a^2} - \frac{y^2}{b^2} = 1$$

where (x, y) are the coordinates of a point on the hyperbola, a is the distance between the hyperbola centre and the intersection H of the hyperbola with the axis connecting the two foci, and b is the length of the line segment, perpendicular to the axis connecting the foci, that extends from H to the asymptote.

The hyperbola in Figure 3 is centred at O , and the distance between O and either focus is denoted by c . Furthermore, it can be shown that $c^2 = a^2 + b^2$ (Kendig, 1938). From the figure, we see that the bearing of the

asymptote is $\beta = \tan^{-1}\left(\frac{b}{a}\right)$. Therefore, in order to solve for β , we must determine the values of b and a .

Figure 3 The *t-range* defines a hyperbola that intersects node R , and whose asymptote passes through the midpoint of the two transmitters in the array (see online version for colours)



We can solve for a by observing that

$$d_{A_1R} - d_{PR} = d_{A_1H} - d_{PH}$$

because, by definition, the distance differences between the foci and all points on the hyperbola are constant. From Figure 3, we see that we can substitute $(c + a)$ for d_{A_1H} and $(c - a)$ for d_{PH} , and therefore

$$d_{A_1R} - d_{PR} = (c + a) - (c - a) = 2a$$

From equation (2), we know the value of $d_{A_1R} - d_{PR}$, which

is the *t-range*, and therefore $a = \frac{d_{A_1PR}}{2}$. We can then solve

for b , using $b = \sqrt{c^2 - a^2}$. In terms of known distances, the bearing of the asymptote is defined as

$$\beta = \tan^{-1} \left(\frac{\sqrt{\left(\frac{d_{A_1P}}{2}\right)^2 - \left(\frac{d_{A_1PR}}{2}\right)^2}}{\left(\frac{d_{A_1PR}}{2}\right)} \right) \quad (3)$$

$$= \cos^{-1} \left(\frac{d_{A_1PR}}{d_{A_1P}} \right).$$

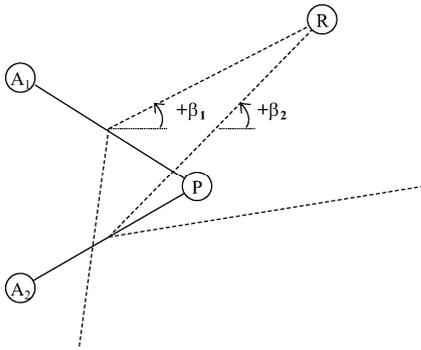
In Figure 3, we see the case where $d_{A_1PR} > 0$, and the position of R lies on the right arm of the hyperbola. If the phase difference is negative (i.e. $\phi_R < \phi_{A_2}$) then the position of R will lie on the left arm of the hyperbola. When this is the case, β is taken clockwise, and we must adjust it by subtracting it from π .

The line $\overline{A_1P}$ connecting the two foci is called the transverse axis of the hyperbola, and is a line of symmetry. This implies that although we know the value of b , we do not know its sign, because mirrored positions on either side of the

transverse axis will result in the same $d_{A,PR}$. Therefore, the asymptote bearing β we obtain using this method could be either positive or negative. To find which bearing is correct, we can switch the roles of the assistant nodes in the array and perform another RIM. This will generate a different t -range, and hence another hyperbolic arm with foci A_2 and P .

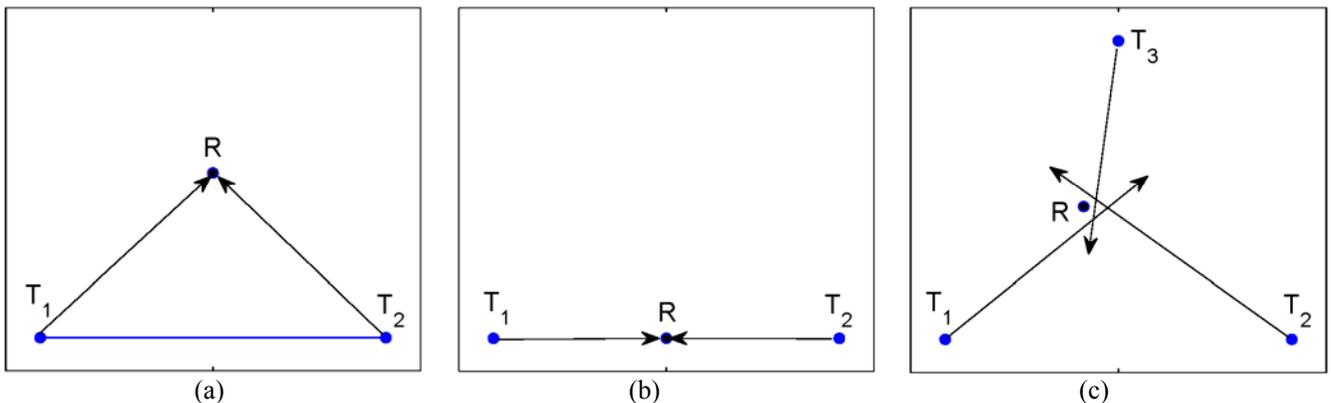
Each hyperbola provides us with two bearing candidates $\pm\beta_i$, where β_i is the angle of the asymptote with the transverse axis, $\overline{A_iP}$. Of course, these angles will be offset from the global x -axis, because the orientation of $\overline{A_iP}$ may not be 0. Adjusting for this, one of the β_1 bearings, and one of the β_2 bearings will point in the same direction, which will approximate the actual bearing of R , as illustrated in Figure 4. Due to the position difference between the centres of the two hyperbolas, we do not expect these two angles to be equal, therefore we define a small threshold ε_β , such that if $|\beta_1 - \beta_2| < \varepsilon_\beta$, these two angles are considered a match. We then take the average of the two angles to obtain our bearing estimate, $\hat{\beta}$.

Figure 4 Determining the true bearing of R is accomplished by selecting $+\beta$ or $-\beta$ from each primary-assistant pair, such that the difference between the two angles is below the threshold ε_β



Because points on the hyperbola converge with the asymptote as their distance from the hyperbola centre increases, the bearing approximation error is larger when R is close to the array. We therefore make the assumption that node R is a sufficient distance from the array. In Section 4,

Figure 5 Triangulation. (a) As few as two bearings from known positions are required to estimate the position of a target. (b) Degenerate case where a third bearing is needed to disambiguate position. (c) Three bearings may not intersect at the same position (see online version for colours)



we show that this distance does not need to be very large when using small-aperture arrays.

3.3 Triangulation

Having approximated the bearing of target node R from a sufficient number of arrays, we can estimate its position using triangulation. Triangulation is the process of determining the position of an object by using the bearings from known reference positions (Esteves et al., 2003). When two reference points are used (Figure 5a), the target position will be identified as the third point in a triangle of two known angles (the bearings from each reference point), and the length of one side (the distance between reference points).

The intersection of bearings can be calculated using the following equations:

$$x = x_2 + \cos(\alpha_2) \frac{y_2 - y_1 - \tan(\alpha_1)(x_2 - x_1)}{\cos(\alpha_2)\tan(\alpha_1) - \sin(\alpha_2)} \quad (4)$$

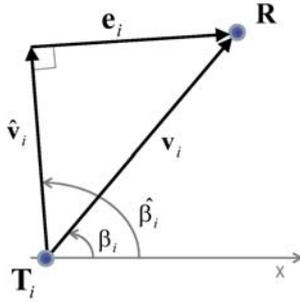
$$y = y_2 + \sin(\alpha_2) \frac{y_2 - y_1 - \tan(\alpha_1)(x_2 - x_1)}{\cos(\alpha_2)\tan(\alpha_1) - \sin(\alpha_2)} \quad (5)$$

where (x, y) are the coordinates of the intersecting bearings (i.e. the position estimate of node R), (x_i, y_i) are the coordinates of the centre of array T_i , and α_i is the bearing of R from T_i .

When the target position is directly between the two reference points (Figure 5b), two bearings are not sufficient to determine position, because the target could be located at any point on that axis. Therefore, a third bearing is required to disambiguate. However, three bearings may not intersect at the same point if any one bearing is inaccurate (Figure 5c). Triangulation techniques are presented in the works of McGillem and Rappaport (1989), Niculescu and Nath (2003), Nasipuri and el Najjar (2006), and Betke and Gurvits (1997), in which position estimation using more than two bearings is considered. The method we chose was a least squares orthogonal error vector solution based on the work of Ash and Potter (2007) and Doğançay (2005), which is rapid, has low complexity, and still provides accurate position estimates from noisy bearing measurements.

Least squares triangulation using orthogonal error vectors works as follows. Figure 6 illustrates a simplified set-up with a single TripLoc anchor (T_i) and target receiver node (R). The actual bearing from the anchor to the target is denoted as β_i and the estimate as $\hat{\beta}_i$. Similarly, the vector pointing from the anchor position to the actual target position is denoted as \mathbf{v}_i and the estimated bearing vector as $\hat{\mathbf{v}}_i$. Note that we have not estimated the length of $\hat{\mathbf{v}}_i$, but only its bearing. Finally, we denote the difference between the actual and estimated bearing vectors as the orthogonal error vector \mathbf{e}_i , such that $\mathbf{e}_i^T \hat{\mathbf{v}}_i = 0$.

Figure 6 Least squares triangulation using orthogonal error vectors (see online version for colours)



If we let $\mathbf{a}_i = \begin{bmatrix} \sin \hat{\beta}_i \\ -\cos \hat{\beta}_i \end{bmatrix}$, then the orthogonal error vector is formally defined as

$$\mathbf{e}_i = \|\mathbf{R} - \mathbf{T}_i\| \sin(\hat{\beta}_i - \beta_i) \mathbf{a}_i,$$

where $\|\mathbf{R} - \mathbf{T}_i\|$ is the distance between the target and anchor position vectors, $\hat{\beta}_i - \beta_i$ is the Gaussian bearing noise with zero mean and variance $\sigma^2_{\beta_i}$, and \mathbf{a}_i is the unit vector orthogonal to $\hat{\mathbf{v}}_i$.

The position of the target can be represented as $\mathbf{R} = \mathbf{T}_i + \hat{\mathbf{v}}_i + \mathbf{e}_i$. To remove $\hat{\mathbf{v}}_i$, we multiply with the transpose of \mathbf{a}_i , resulting in

$$\mathbf{a}_i^T \mathbf{R} = \mathbf{a}_i^T \mathbf{T}_i + \eta_i,$$

where $\eta_i = \|\mathbf{R} - \mathbf{T}_i\| \sin(\hat{\beta}_i - \beta_i)$. Considering all anchors ($i = 1, \dots, N$), we have a system of equations that take the form

$$\mathbf{A}\mathbf{R} = \mathbf{b} + \boldsymbol{\eta}$$

where $\mathbf{A} = [\mathbf{a}_1^T, \mathbf{a}_2^T, \dots, \mathbf{a}_N^T]^T$ and $\mathbf{b} = [\mathbf{a}_1^T \mathbf{T}_1, \mathbf{a}_2^T \mathbf{T}_2, \dots, \mathbf{a}_N^T \mathbf{T}_N]^T$. A least squares solution for estimating \mathbf{R} is given by

$$\hat{\mathbf{R}} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$

where $\hat{\mathbf{R}}$ is the position estimate returned by the triangulation using noisy bearing measurements from N anchors.

4 Error analysis

In this section, we present an error analysis of our bearing and position estimation techniques using TripLoc. It is important to note that, although we use phase differences as input to our bearing estimation algorithm, the algorithm is generalisable to small-aperture sensor arrays that can derive distance differences using *any* means. For instance, RF ultra wide band antenna arrays, acoustic or ultrasonic sensors and other types of arrays that can yield time-difference-of-arrival measurements from (sufficiently) distant sources fall into this category. Therefore, in this section, we assume the inputs to be distance differences. Notice that the distance differences are linearly related to RIM measurements [see equation (2)], and therefore the error sensitivity results presented below remain valid. In the generalised case, the same applies to TDOA measurements, from which the distance differences can be computed via multiplication of the respective signal propagation speed (speed of sound for acoustic, speed of light for RF).

4.1 Bearing estimation error analysis

Typically, bearings are computed from distance differences by solving a nonlinear set of equations using iterative techniques. Such techniques are prohibitive on low-end microprocessors due to their computational cost. We make a set of assumptions that allows us to compute bearing estimates in a reasonably simple way. While our bearing estimation technique is computationally less expensive than traditional nonlinear optimisation techniques, our simplifying assumptions introduce estimation errors, which we identify below.

- *Measurement noise*: The distance differences observed by the receiver nodes contain measurement noise. The measurement noise can be attributed to, for instance, non-ideal signal propagation, noise from the electrical circuitry of the receiver, sampling error and quantisation error of the analogue-to-digital converter.
- *Asymptote approximation*: For a pair of transmitters, we approximate the bearing of the receiver with the angles of the asymptotes of the hyperbola. This is a good approximation if the receiver is sufficiently far from the transmitters, because the hyperbola converges on its asymptote. However, for close receivers, errors due to this assumption will not be negligible.
- *Translation of bearing candidates*: At least two transmitter pairs are required to unambiguously compute the bearing because, for just one transmitter pair, the angles of both asymptotes are possible

solutions. Since, for a transmitter pair, we compute the bearing candidates with respect to the midpoint of the segment defined by the two antennas, fusing bearing candidates from two different transmitter pairs is not possible without knowing the distance of the receiver. We use the far-field assumption (i.e. that the receiver is infinitely far from the transmitter array) to carry out the disambiguation and fusion of bearings, introducing an error this way.

We intentionally omit the analysis of array position and orientation errors and instead make the following assumptions:

- *Antenna configuration is known:* The transmitter locations are assumed to be given. It is assumed that the transmitter nodes are fabricated with a prescribed antenna separation.
- *Relative bearings:* We assume that the computed bearings are given in the local coordinate system of the array. Hence, the location and orientation errors of the array are not considered in the error analysis of the bearing estimation.

We first analyse the sensitivity of the bearing estimates to noise in the distance difference inputs. Second, we analytically derive the errors related to the asymptote approximation and to the translation of bearing candidates. These errors depend on the bearing and distance of the target receiver, relative to the transmitter array. Finally, we provide an analysis of the total bearing estimation error resulting from both noise in the inputs and the errors due to the asymptote approximation and the translation of bearing candidates.

4.1.1 Sensitivity of bearing to measurement noise

A distance difference from a pair of transmitters in the array constrains the location of the receiver to one arm of a hyperbola, the foci of which are the positions of the two

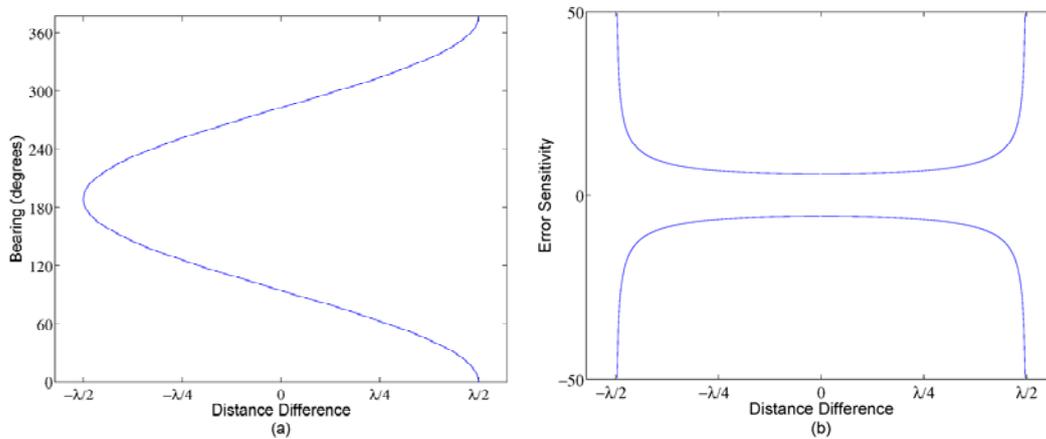
transmitters. For the sake of simplicity, let us assume that the two transmitters P and A_1 are located at $(c, 0)$ and $(-c, 0)$, respectively (see Figure 3). If the measured distance difference is positive, the receiver is constrained to the right arm of the hyperbola, while if the distance difference is negative, the receiver is located on the left arm. We approximate the bearing of the receiver using the asymptote angles as follows:

$$\hat{\beta} = \begin{cases} \pm \tan^{-1} \left(\sqrt{\frac{c^2 - a^2}{a}} \right) & \text{if } a > 0 \\ \pm \frac{\pi}{2} & \text{if } a = 0 \\ \pi \mp \tan^{-1} \left(\sqrt{\frac{c^2 - a^2}{a}} \right) & \text{if } a < 0 \end{cases} \quad (6)$$

We analyse the sensitivity of the bearing estimates $\hat{\beta}$ to noise in the distance difference by taking the partial derivative of equation (6) with respect to the distance difference. To see what amplification effect an error in a given distance difference d produces on the bearing estimate, we need to evaluate the partial derivative at d .

Figure 7a shows the relation between the measured distance difference d and the bearing candidates $\hat{\beta}$ when the antenna separation is half the wavelength ($\lambda/2$). Notice the ambiguity of the bearing candidates. Figure 7b plots $\delta\hat{\beta}/\delta d$ for each solution of $\hat{\beta}$. This figure shows that when the absolute value of the measured distance difference is close to the antenna separation, the computed bearing candidates are very sensitive to measurement noise. For instance, if the distance difference is 80% of the antenna separation, an infinitesimally small error in the measurement will be amplified tenfold in the bearing estimate.

Figure 7 (a) Relationship between measured distance difference and computed bearing, (b) sensitivity of the computed bearing to measurement noise (see online version for colours)



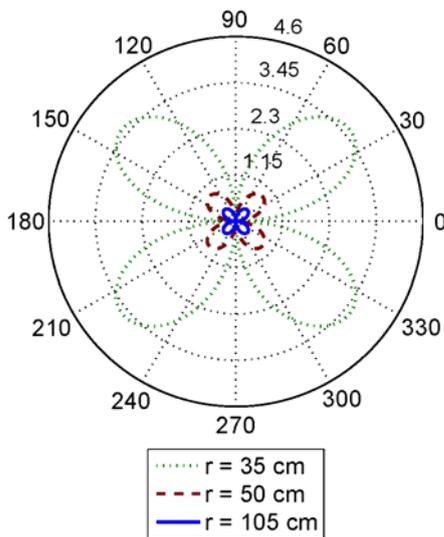
4.1.2 Asymptote approximation

The error of approximating a hyperbola with its asymptote is the difference between the approximated bearing $\hat{\beta}$ and the actual bearing β of the receiver. Assuming that the receiver R is located at (u, v) , the actual bearing will be $\beta = \tan^{-1}(u/v)$. Hence, the error $\epsilon = \hat{\beta} - \beta$ introduced by the asymptote assumption is

$$\epsilon = \begin{cases} \pm \tan^{-1}\left(\frac{v}{u}\right) \mp \tan^{-1}\left(\frac{\sqrt{c^2 - a^2}}{a}\right) & \text{if } a > 0 \\ 0 & \text{if } a = 0 \\ \mp \tan^{-1}\left(\frac{v}{u}\right) \mp \tan^{-1}\left(\frac{\sqrt{c^2 - a^2}}{a}\right) & \text{if } a < 0 \end{cases} \quad (7)$$

Figure 8 shows the error introduced by the asymptote approximation when the receiver is located respectively one, two and three times the antenna distance away from the midpoint of the segment connecting the two antennas. As expected, the error of the approximation decreases as the distance of the receiver from the transmitter array increases, that is, as the hyperbola converges on its asymptote. As we can see, the maximum error introduced by the asymptote assumption is less than 0.6° , as little as three antenna distances away.

Figure 8 Error in bearing (in degrees) caused by the assumption that the receiver lies on the hyperbola asymptote (see online version for colours)



4.1.3 Translation of bearing candidates

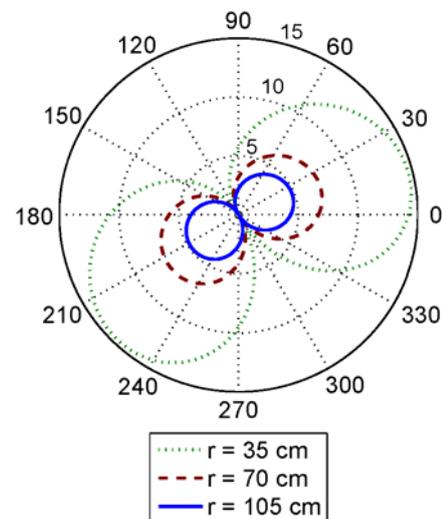
For a pair of transmitter antennas, it is not possible to unambiguously approximate the bearing of the asymptote. Because the hyperbola arm has two asymptotes, the angle of either one can be the correct bearing estimate. Hence, we need two transmitter antenna pairs for disambiguation. Let us treat the bearing candidates (computed from the t -ranges

of both transmitter antenna pairs) as vectors of unit length, with bases at the centre of the hyperbolas and whose angles are the computed bearing candidates. Since these vectors are given in the coordinate system of the respective hyperbolas, we need to transform them to the coordinate system of the array. This transformation includes a translation and a rotation. Then, we translate each vector such that its base is at the origin. Clearly, the bearing vector translated this way will not point directly toward the target receiver anymore, but if the receiver is sufficiently far from the transmitter array, the introduced angular error will be small. Finally, we disambiguate the bearing candidates by finding two that have approximately the same value.

Let us now express the angular error caused by the translation of bearing candidates. We assume that the transmitter is a uniform circular array of three antennas, with a pairwise antenna distance of less than $\lambda/2$. The coordinate system of the array is set up such that the midpoint of the array is at the origin, and antenna P lies on the positive side of the x -axis. Let us consider only the correct bearing candidate (the other will be discarded later) for the transmitter pair P and A_1 . Furthermore, let us assume for now that the bearing candidate has no error. The difference between the actual bearing of the target receiver and the angle of the bearing candidate translated to the origin gives the angular error of the far-field assumption.

Figure 9 shows the error introduced by the far-field assumption when the receiver is located respectively one, two and three times the antenna distance away from the midpoint of the segment connecting the two antennas. As we can see, as few as three antenna distances away, the maximum error introduced by the far-field assumption is less than 5° . In this particular antenna arrangement, the maximum errors are at 15° and 225° , respectively.

Figure 9 Error in bearing (in degrees) caused by the assumption that the bearing from the midpoint of the segment connecting the two antennas equals the bearing from the origin of the array coordinate system (see online version for colours)



4.1.4 Compound bearing estimation error

Since one transmitter pair reports two bearing candidates, at least two transmitter pairs are required to resolve this ambiguity. For the sake of simplicity, let us assume that we have two transmitter pairs. Clearly, there must be one bearing candidate for each transmitter pair that is close to the true bearing. Except for some degenerate cases, the other two bearing candidates will be significantly different than the true bearing and will not be close to each other (see Figure 4). Therefore, in order to disambiguate between the correct and incorrect bearing candidates, we take all possible pairs of bearing candidates, one from the first transmitter pair and the other from the second transmitter pair and find the pair with the least pairwise angular difference. The reported bearing estimate is computed as the average of the two closest bearing candidates.

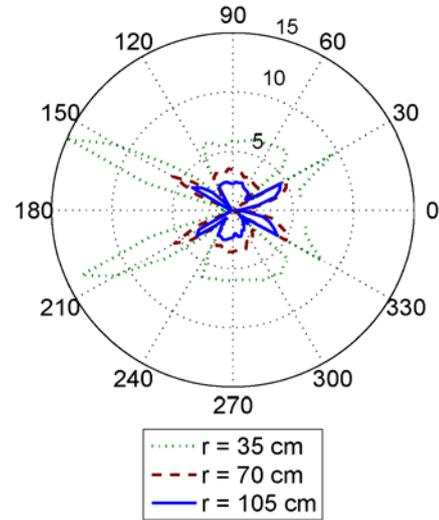
Figure 10 shows the bearing estimation errors considering the above three types of error sources, averaged over 500 simulation rounds. We added a Gaussian noise to the distance differences, with zero mean and standard deviation set to 5% of the antenna distance. The plot suggests that the expected bearing estimation errors are below 5°, and peak around 30°, 150°, 240° and 330°, exactly where the individual transmitter pairs exhibit high error sensitivity.

4.2 Position estimation error analysis

Our localisation technique uses triangulation to compute the position of the receiver nodes. Triangulation is a well-studied localisation method (McGille and Rappaport, 1989; Betke and Gurvits, 1997; Niculescu and Nath, 2003; Nasipuri and el Najjar, 2006); therefore, we omit a detailed evaluation and refer the reader to the above references. Instead, we only present simulation results that demonstrate the effects of error from different sources. The simulation results will also serve as a reference to evaluate our experimental results.

We consider three different sources of error: position error of the arrays, orientation error of the arrays and bearing estimation error.

Figure 10 Absolute error of bearing estimation (in degrees) caused by noisy distance differences, averaged over 500 simulation rounds. The standard deviation of the distance difference errors is 5% of the antenna distance (see online version for colours)

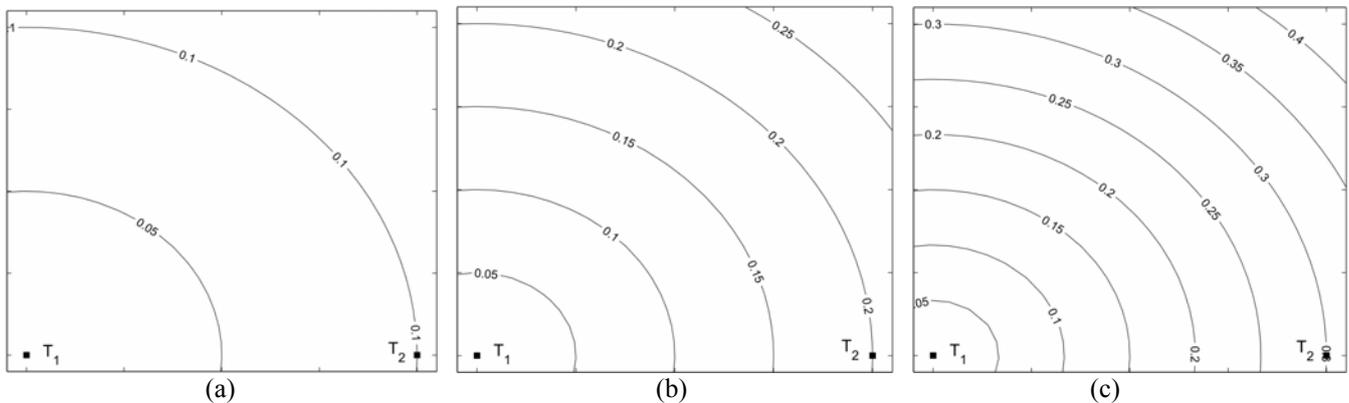


4.2.1 Position error of the arrays

For simplicity, let us assume that we have two arrays T_1 and T_2 , located at $(0, 0)$ and $(u, 0)$, respectively. We model the error of array positioning by varying u around a nominal value of 20 m. Note that we can omit analysis of the effect of varying the y -coordinate of T_2 , since an error in both coordinates of T_2 can be modelled by an error in the x -coordinate and an orientation error of T_1 .

Figure 11 shows the localisation error caused by error in the x -coordinate of array T_2 , while all other sources of error are eliminated. As we can see, the maximum errors in the $20\text{ m} \times 20\text{ m}$ sensing region are approximately 0.15 m, 0.3 m and 0.5 m for array position errors of 0.1 m, 0.2 m and 0.3 m, respectively.

Figure 11 Contour map of location errors (in metres) caused by (a) 0.1 m, (b) 0.2 m and (c) 0.3 m error in the x coordinate of array T_2 , respectively (see online version for colours)



4.2.2 Orientation error of the arrays

To simulate the effects of orientation error, we vary the orientation of array T_2 around its nominal value. This is simulated by adding a constant to the bearings reported from T_2 .

Figure 12 presents the localisation error due to 1, 4 and 8° of orientation error of T_2 , eliminating all other error factors. As expected, if the receiver is close to the line defined by the two arrays, even a small orientation error will cause a considerable localisation error. This is because the lines defined by the bearings intersect at an angle close to 180°, and even a small error in the angles causes a considerable relocation of the intersection. The simulation results show that this particular configuration is very sensitive to orientation errors: with an orientation error of 1°, we see 0.5 m localisation errors in the sensing region. With an orientation error of 4°, the expected errors are close to 3 m, while with 8° orientation error, the localisation errors become unacceptably large in most of parts of the sensing region.

4.2.3 Bearing estimation error

We now present simulation results to characterise the effect of bearing estimation errors on the computed locations. We

assume that the location and the orientation of two antenna arrays are known. Furthermore, we assume that the bearing estimates have additive white noise. We model the noise with a Gaussian random variable with zero mean and a known standard deviation σ . It is assumed that the noise in the bearings from arrays T_1 and T_2 are uncorrelated.

Through 100 simulation rounds, we compute the standard deviation of localisation errors for all points on a high-resolution grid covering the entire sensing region. Figure 13 shows the simulation results for bearing errors with standard deviation $\sigma = 4^\circ$, $\sigma = 5^\circ$ and $\sigma = 6^\circ$, respectively. The results suggest that standard deviation of the position errors is high when the receiver is close to the line connecting the two arrays, and increases towards the edge of the sensing region. The central part of the sensing region gives the lowest error variance. As we increase the standard deviation of the bearing error, the standard deviation of the location error also increases. As the standard deviation of the bearing error increases, localisation precision degrades. With 6° standard deviation of bearings, only the central part of the sensing region offers less than 1.5 m location error variance.

Figure 12 Contour map of location errors (in metres) caused by (a) 1°, (b) 4° and (c) 8° errors in the orientation of array T_2 , respectively

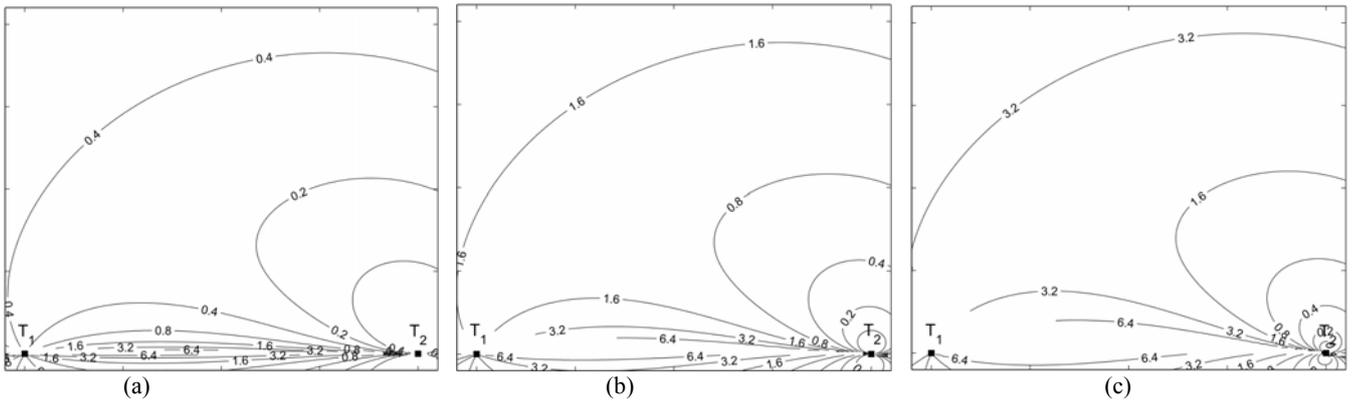
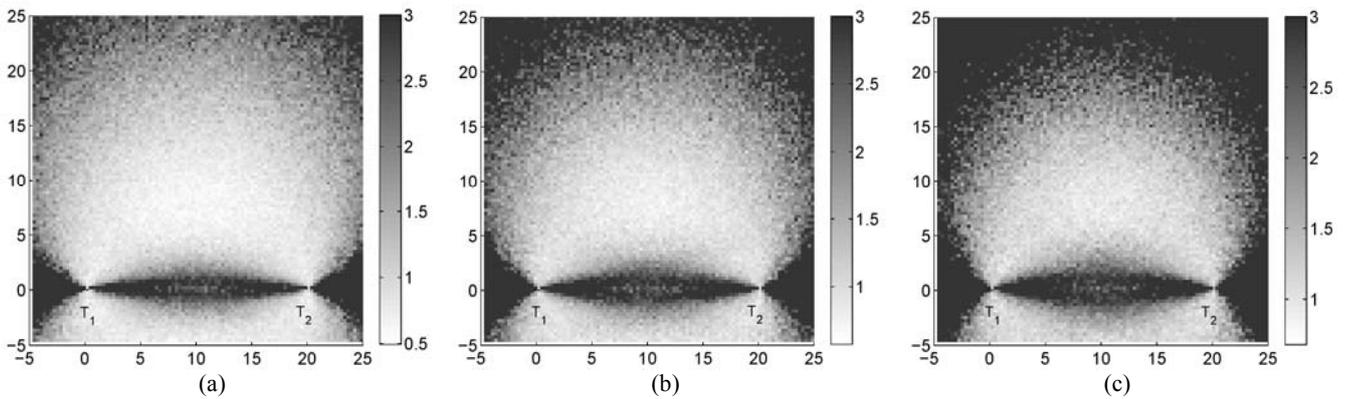


Figure 13 Standard deviation of location error (in metres) caused by bearing noise with zero mean and standard deviation (a) $\sigma = 4^\circ$, (b) $\sigma = 5^\circ$ and (c) $\sigma = 6^\circ$, respectively



5 Implementation

Our system is implemented on Crossbow ExScal motes (Dutta et al., 2005), which use the Texas Instruments CC1000 radio chip (Texas Instruments, 2007), and transmit in the 433 MHz range. Three XSMs form a TripLoc array. Because the two transmitting antennas are so close to each other, they will suffer from parasitic effects (Carr, 2001). To minimise this negative interference, we place the nodes in a mutually orthogonal configuration. The nodes are elevated approximately 1.5 m to reduce ground-based reflections. The TripLoc antenna array is pictured in Figure 1.

All nodes in our system execute the same distributed application, coded in nesC (Gay et al., 2003), and run the TinyOS operating system (Hill et al., 2000). All operations run locally, and there is no offline or PC-based processing. The entire application requires 3 kB of RAM and 57 kB of program memory (ROM).

5.1 Antenna separation

The antenna separation of the TripLoc array is 35.35 cm. In order to avoid the modulo 2π phase ambiguity, we were required to transmit a signal with wavelength greater than 70.7 cm (i.e. a frequency less than 424 MHz). This was possible, as we observed reliable CC1000 transmissions as low as 405 MHz. Ultimately, we chose to transmit at a carrier frequency of 409 MHz, corresponding to a wavelength of 73 cm.

5.2 Initialisation

Before starting the localisation process, system initialisation is required.

5.2.1 Scheduling

First, we must create a schedule declaring the order that array anchors perform their RIMs. This schedule is broadcast to all anchors in the sensing region. During runtime, the anchors keep track of who is currently performing the RIM. The next anchor in the schedule will then know to begin its RIM when the previous one ends.

5.2.2 Tuning

One of the issues that arises when using the CC1000 is that the true carrier frequency can differ from that of the desired frequency by up to 2 kHz. In order to account for this, we must tune the radios by having the primary and assistant nodes transmit at the intended frequencies, while a third node monitors the signal. The primary node then changes its transmission frequency in small steps. The receiver observes the resulting beat frequency, and when it matches the desired one, informs the primary node. From this point on, the primary node transmits at this new frequency, in order to maintain a consistent interference signal.

5.2.3 Calibration

Finally, we note that it is difficult to fix exact array orientations. If the array is not level, or is rotated slightly, the shape of the resulting hyperboloid will degrade the localisation result. An

easy way around this is to place one or more sensor nodes at known positions and run the bearing estimation algorithm. Because all positions are known, bearings between the sensor nodes and arrays are also known. The difference between the actual bearings and the computed bearings indicate the true orientation of the arrays. By using this simple, yet effective technique, we saw the bearing error decrease by 56%.

5.3 Radio interferometric measurements

We perform radio interferometric measurements using the technique presented and evaluated by Maróti et al. (2005). A RIM consists of several tasks, executed sequentially according to a strict schedule. In order to transmit or receive the RF signal at the desired frequency, the radio must be acquired from the MAC layer. Once the radio is acquired, nodes are unable to use their radios for standard message passing, so each node participating in the RIM executes each task in lockstep. Therefore, before acquiring the radio from the MAC layer, the first step is to synchronise the participating nodes.

5.3.1 Time synchronisation

Because phase difference is used to determine bearing, each node must measure the signal phase at the same time instant. This requires synchronisation with accuracy on the order of microseconds or better. To accomplish this, we use a SyncEvent (Kusý et al., 2006), the same technique employed by RIPS and other radio interferometric based ranging methods. A message is broadcast by the primary transmitter that contains a time in the near future to start the RIM. The time is specified according to the local clock of the primary node. The clock offset between the sender and receiver of this message is determined as follows. The sender inserts a timestamp at the end of the message *after* the message has already begun transmission. On the receiver side, a timestamp is taken upon message reception. Because the propagation time of the message through air is negligible, the difference between these two timestamps approximates the clock offset between the two nodes. Clock skew is not an issue here because synchronisation is performed at the beginning of the RIM, and relative clock rate does not noticeably change in the time it takes to perform the RIM. One advantage of this type of synchronisation is that any number of receiver nodes within the broadcast range can participate, because the broadcast message is received at all nodes at approximately the same time.

5.3.2 Acquiring the radio

The first task each node (transmitter and receiver) must perform after synchronisation is acquiring the radio from the MAC layer. This involves switching radio drivers to enable the transmission/reception of a pure sine wave at the desired frequency, a feature the standard radio driver does not support.

5.3.3 Radio calibration

Once the radio has been acquired, we must calibrate it for transmission at the desired frequency. To do this, we use previous calibration data (as described by Maróti et al., 2005), which avoids the time-consuming calibration step that automatically occurs in the restore task.

5.3.4 Signal transmission/reception

In order to determine the signal phase, we must sample enough RSSI values to reconstruct the signal. This requires a transmission/reception time that is long enough to sample about six periods, from which we can use filtering techniques to obtain a fairly smooth sinusoid. As each RSSI value is obtained, it is inserted into a 256-byte buffer. At the end of the sampling period, we traverse the buffer, keeping track of the minimum and maximum values, from which we can determine the signal amplitude and frequency. Once these signal characteristics are known, we can determine the phase of the signal at the specified time instant.

5.3.5 Restoring the radio

Once the transmission is complete, the standard radio must then be restored to its original state in order to exchange the phase data between nodes.

5.4 Bearing approximation and triangulation

After each array performs its RIMs, the assistant nodes broadcast their observed phase data, which are received by the target receiver nodes. The phase data packets are small, only 12 bytes long, five of which comprise the standard active message header, two for the assistant ID, one for the sequence number and four bytes for the measured phase. For redundancy, the data packets are broadcast three times (with a small random delay to avoid contention between nodes); however, any messages that arrive after the communication deadline of 250 ms will be discarded. The communication deadline signals the start of the localisation process. The localisation method performs both bearing approximation and triangulation and is optimised for memory and speed. The localisation result is sent to a base station for logging; however, for mote applications that require the position updates, no routing is required and the position estimate is readily available.

5.5 Latency analysis

Because we would like to use TripLoc to localise mobile sensors in addition to stationary nodes, each array must perform its RIMs as fast as possible so that the mobile sensor has not had a chance to significantly change its position. In order to keep the localisation latency to a minimum, we provide an analysis of the different component execution times.

Table 1 lists the execution deadlines of the tasks presented above. These deadlines were chosen to give each task enough time to complete, assuming a reasonable amount of jitter and were determined experimentally. Each array performs two RIMs, one for each primary-assistant pair. For each array, only one synchronisation message is broadcast at the beginning of the first RIM. When a RIM completes, a *measurement Ended()* event is triggered on each participating node. This includes nodes belonging to other arrays, even if they took no part in the phase measurement. The event enables the next array in the schedule to begin its RIM without delay. The receiver nodes

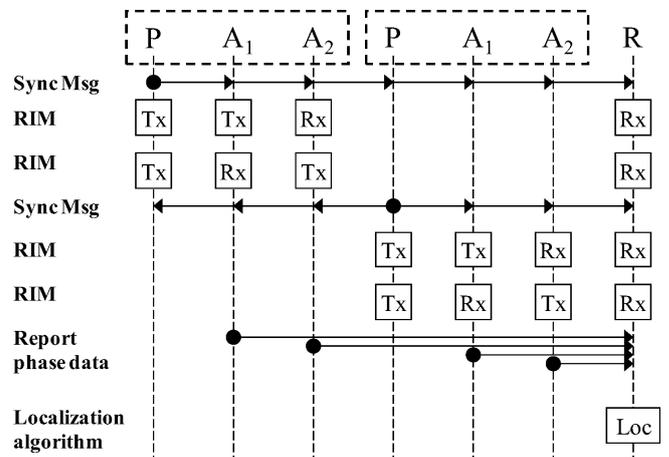
store their phase measurements for each array until all arrays have finished their RIMs. After all arrays have finished, the participating assistants broadcast their phase measurements and the receiver nodes calculate their bearings to each array. We allow 125 ms for sending phase measurements. Localisation takes no longer than 6.5 ms, after which the first array begins its next round of RIMs.

Table 1 Latency of localisation tasks

Task	Latency (ms)
Synchronise	163
Acquire	5.4
Calibrate	1.08
Transmit/Receive	63.2
Restore	49.91
Report phase	125
Localise	6.5

Figure 14 is a sequence diagram of these events using a set-up of two arrays and a single receiver node. Because the target nodes act as receivers, no additional latency is incurred by introducing more to the sensing region. Each array sends one synchronisation message and performs two RIMs, for a total time of 402.18 ms. After all arrays have performed their RIMs, an additional 130.5 ms is required for communication and localisation. Therefore, localisation using two arrays will take no longer than 0.93 s, three arrays will take no longer than 1.34 s and no longer than 1.74 s when using four arrays. It should be noted that this latency can be further decreased by removing subsequent synchronisation messages after the first one. Because a round of RIMs takes a relatively short time, clocks will not experience a significant amount of drift. In addition, the time allotted to report phases is much greater than is actually needed and can be shortened, if necessary. We chose to set it longer to ensure we received all measurement data for analysis and debugging purposes. From this latency analysis, we see that a sufficient number of arrays can perform RIMs to enable a receiver node to determine its position in less than 1 s.

Figure 14 Sequence diagram of a RIM schedule with two arrays (dotted boxes) and a receiver node (R)



6 Evaluation

To evaluate the accuracy of TripLoc, we performed three sets of experiments. The first two experiments measure the bearing accuracy of the system and the third experiment measures the localisation accuracy.

6.1 Bearing accuracy

For Experiment 1, we measure the bearing accuracy of six receiver nodes, which are evenly spaced around the array every 60° at a distance of 10 m from the array centre (Figure 15a). This experiment demonstrates how the bearing error changes with respect to array orientation. We performed 50 bearing estimates for each node surrounding the array. The average bearing errors are displayed in Figure 16.

For Experiment 2, we measure the bearing accuracy of 14 receiver nodes from three arrays surrounding the sensing region in an outdoor, low-multipath environment (Figure 15b). This experiment is more representative of a real-world deployment with multiple anchors. We performed approximately 35 bearing estimates from each anchor to all target nodes, resulting in a total of 105 estimates per target and 1470 estimates total. The distribution of bearing estimate errors is shown in Figure 17. The average bearing estimation error is 3.2° overall, with a 6.4° accuracy at the 90th percentile. The errors from both experiments are consistent with our bearing error analysis in Section 4.

Figure 15 Experimental set-up for bearing accuracy of TripLoc: (a) Experiment 1: Bearing accuracy of one array. Six receiver nodes ($R_1 \dots R_6$) are placed 10 m from array (A), with angular separation of 60° , (b) Experiment 2: Three arrays ($A_1 \dots A_3$) surround the 20×20 m sensing region containing 14 target receiver nodes ($R_1 \dots R_{14}$) (see online version for colours)

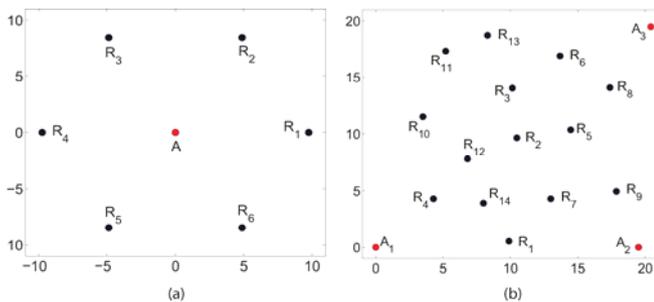


Figure 16 Experimental results: Experiment 1 average bearing error with respect to array orientation (sample size of 50)

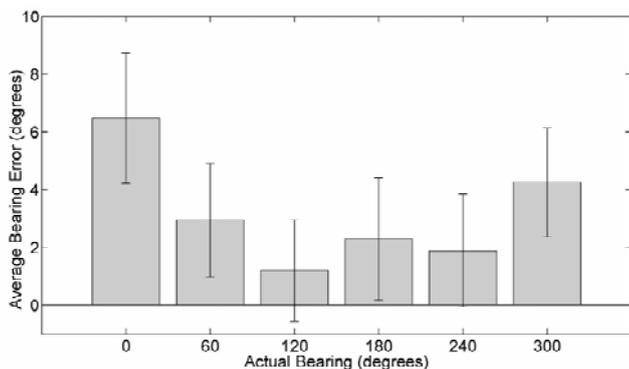
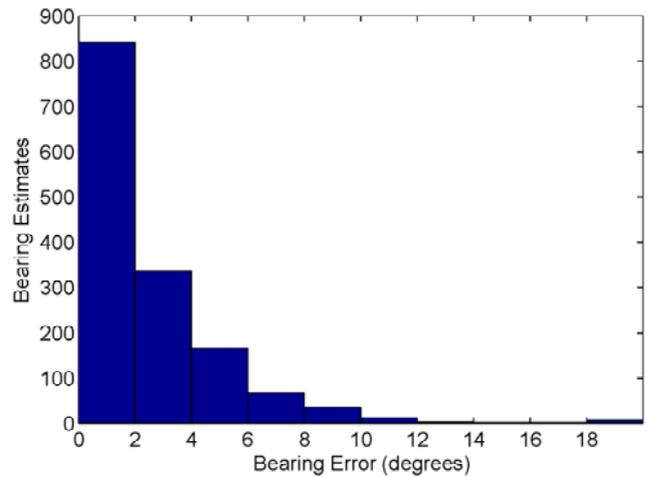


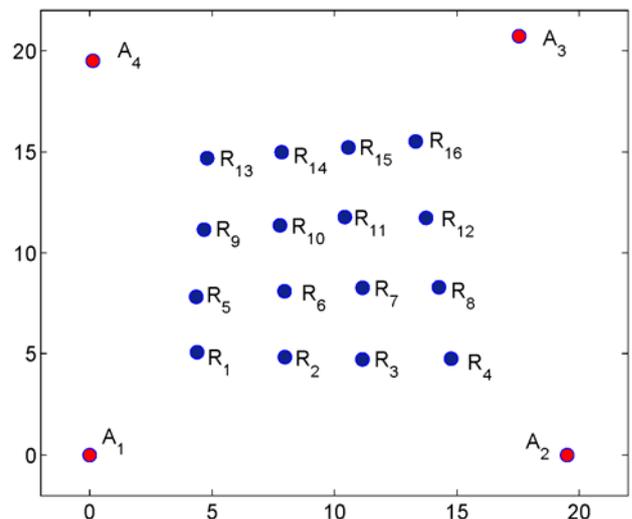
Figure 17 Experimental results. Experiment 2 bearing error distribution (sample size of 1470) (see online version for colours)



6.2 Localisation accuracy

For Experiment 3, we place four TripLoc array anchors at the corners of a 20×20 m sensing region containing 16 target receiver nodes in a slightly randomised grid. The anchor nodes are set back from the target nodes in order to minimise the position error that results at sensitive bearings (see Section 4; an evaluation of node placement at sensitive bearings was performed by Amundson, 2010). Each target node periodically estimates its position using the bearing estimation and triangulation techniques presented in Section 3. Figure 18 illustrates our set-up for this experiment.

Figure 18 Experimental set-up for localisation accuracy of TripLoc (Experiment 3). Four arrays ($A_1 \dots A_4$) surround the 20×20 m sensing region containing 16 target receiver nodes ($R_1 \dots R_{16}$) (see online version for colours)



For Experiment 3, we perform 100 position estimates for each target node. Figure 19 shows the average position estimate for each target relative to its actual position. The distribution of position estimate errors is shown in Figure 20. The average overall position error is 0.78 m.

Figure 19 Experiment 3 average position estimate for each target relative to actual position. Arrows connect the estimated position of a target node with its actual location (see online version for colours)

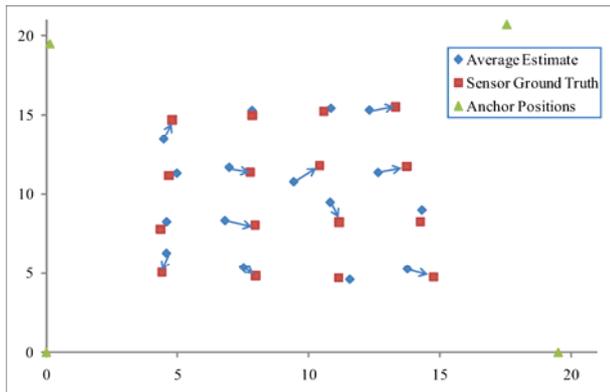
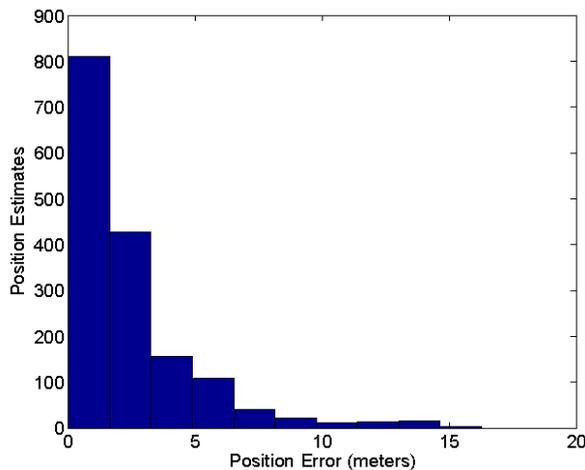


Figure 20 Experiment 3 position error distribution (sample size of 1600) (see online version for colours)



7 Conclusion

In this paper, we present a method for rapid distributed bearing estimation and localisation in WSNs. The array anchor in our system consists of three nodes, two of which transmit at frequencies that interfere to create a low-frequency beat signal. The phase of this signal is measured by the third node in the array, as well as by multiple target nodes at unknown positions. The phase difference defines a hyperbola, and bearing can be approximated by calculating the angle of the asymptote. Our experimental results show that this technique has an average bearing estimation accuracy of 3.2° , average position accuracy of 0.78 m and measurements can be taken in approximately 1 s.

Our system is designed to overcome several challenges in WSN AOA localisation. The array prototype is easily constructed by fixing three nodes together with antennas at orthogonal angles. It is comprised entirely of COTS sensor nodes, and no additional hardware is required because RIM-based ranging only requires use of the mote radio. Unlike other radio interferometric techniques, our system avoids the modulo 2π ambiguity, and therefore the need to perform RIMs on multiple

channels, by separating the two transmitting antennas less than half the wavelength of the carrier frequency. Similarly, by constraining the location of one of the RIM receivers to the array, it becomes possible to approximate the bearing of the other receiver without prolonged computation or having to rely on a base station for processing.

Acknowledgements

This work was supported by ARO MURI grant W911NF-06-1-0076, NSF grant CNS-0721604, NSF CAREER award CNS-0347440 and the TAMOP-4.2.2/08/1/2008-0008 program of the Hungarian National Development Agency. The authors would also like to thank Peter Volgyesi for his assistance, and Metropolitan Nashville Parks and Recreation, and Edwin Warner Park for use of their space.

References

- Altun, K. and Koku, A.B. (2005) 'Evaluation of egocentric navigation methods', *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication (ROMAN)*, August, pp.396–401.
- Amundson, I. (2010) *Spatio-Temporal Awareness in Mobile Wireless Sensor Networks*, PhD Dissertation, Vanderbilt University.
- Amundson, I., Koutsoukos, X. and Sallai, J. (2008) 'Mobile sensor localization and navigation using RF doppler shifts', *Proceedings of the 1st ACM International Workshop on Mobile Entity Localization and Tracking in GPS-less Environments (MELT)*, ACM, San Francisco, CA, USA, pp.97–102.
- Ash, J.N. and Potter, L.C. (2004) 'Sensor network localization via received signal strength measurements with directional antennas', *Proceedings of the Allerton Conference on Communication, Control, Computing*, pp.1861–1870.
- Ash, J.N. and Potter, L.C. (2007) 'Robust system multiangulation using subspace methods', *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN)*, pp.61–68.
- Bahl, P. and Padmanabhan, V.N. (2000) 'Radars: an in-building RF-based user-location and tracking system', *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, March, Vol. 2, 775–784.
- Bekris, K.E., Argyros, A.A. and Kavraki, L.E. (2004) 'Angle-based methods for mobile robot navigation: reaching the entire plane', *Proceedings of the International Conference on Robotics and Automation (ICRA)*, New Orleans, LA, pp.2373–2378.
- Betke, M. and Gurvits, L. (1997) 'Mobile robot localization using landmarks', *IEEE Transactions on Robotics and Automation*, Vol. 13, No. 2, pp.251–263.
- Biswas, P., Aghajan, H. and Ye, Y. (2006) 'Semidefinite programming based algorithms for sensor network localization', *Proceedings of the 39th Asilomar Conference on Signals, Systems and Computers*, pp.188–220.
- Capon, J. (1969) 'High-resolution frequency-wavenumber spectrum analysis', *Proceedings of the IEEE*, Vol. 57, No. 8, ISSN 0018-9219.
- Carr, J. (2001) *Practical Antenna Handbook*, 4th ed., McGraw Hill.

- Chang, H., Tian, J., Lai, T., Chu, H. and Huang, P. (2008) 'Spinning beacons for precise indoor localization', *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys)*, Raleigh, NC, USA, pp.127–140.
- Chang, H., You, C., Chu, H. and Huang, P. (2007) 'Modeling and optimizing positional accuracy based on hyperbolic geometry for the adaptive radio interferometric positioning system', *Proceedings of the 3rd International Workshop on Location and Context Awareness (LoCA)*, Vol. 4718, pp.228–244.
- Doğançay, K. (2005) 'Bearings-only target localization using total least squares', *Signal Processing*, Vol. 85, No. 9, pp.1695–1710.
- Dutta, P., Grimmer, M., Arora, A., Bibyk, S. and Culler, D. (2005) 'Design of a wireless sensor network platform for detecting rare, random, and ephemeral events', *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN/SPOTS)*, IEEE Press Piscataway, NJ, USA, Article No. 70.
- Esteves, J.S., Carvalho, A. and Couto, C. (2003) 'Generalized geometric triangulation algorithm for mobile robot absolute self-localization', *Proceedings of the IEEE International Symposium on Industrial Electronics (ISIE)*, pp.346–351.
- Friedman, J., Charbiwala, Z., Schmid, T., Cho, Y. and Srivastava, M. (2008) 'Angle-of-arrival assisted radio interferometry (ARI) target localization', *Proceedings of the Military Communications Conference (MILCOM)*, pp.1–7.
- Gay, D., Levis, P., von Behren, R., Welsh, M., Brewer, E. and Culler, D. (2003) 'The nesC language: a holistic approach to networked embedded systems', *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pp.1–11.
- Harter, A., Hopper, A., Steggle, P., Ward, A. and Webster, P. (1999) 'The anatomy of a context-aware application', *In Mobile Computing and Networking*, p.5968.
- Hightower, J. and Borriello, G. (2001) 'Location systems for ubiquitous computing', *IEEE Computer*, Vol. 34, No. 8, pp.57–66.
- Hill, J., Szewczyk, R., Woo, A., Hollar, S., Culler, D. and Pister, K. (2000) 'System architecture directions for networked sensors', *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, November, pp.93–104.
- Kendig, K. (1938) *Conics*, Mathematical Association of America.
- Kumaresan, R. and Tufts, D.W. (1983) 'Estimating the angles of arrival of multiple plane waves', *IEEE Transactions on Aerospace and Electronic Systems*, AES, Vol. 19, No. 1, ISSN 0018-9251.
- Kusý, B., Dutta, P., Levis, P., Maróti, M., Lédeczi, A. and Culler, D. (2006) 'Elapsed time on arrival: a simple and versatile primitive for canonical time synchronization services', *International Journal of Ad Hoc and Ubiquitous Computing*, Vol. 2, No. 1, pp.239–251.
- Kusý, B., Lédeczi, A. and Koutsoukos, X. (2007a) 'Tracking mobile nodes using RF doppler shifts', *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems (SenSys)*, Sydney, Australia, pp.29–42.
- Kusý, B., Sallai, J., Balogh, G., Lédeczi, A., Protopopescu, V., Tolliver, J., DeNap, F. and Parang, M. (2007b) 'Radio interferometric tracking of mobile wireless nodes', *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services (MobiSys)*, pp.139–151.
- Lucarelli, D., Saksena, A., Farrell, R. and Wang, I. (2008) 'Distributed inference for network localization using radio interferometric ranging', in Verdone, R. (Ed.): *EWSN*, Vol. 4913 of *Lecture Notes in Computer Science*, Springer, pp.52–73.
- Maróti, M., Kusý, B., Balogh, G., Völgyesi, P., Nádas, A., Molnár, K., Dóra, S. and Lédeczi, A. (2005) 'Radio interferometric geolocation', *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems (SenSys)*, pp.1–12.
- McGillem, C.D. and Rappaport, T.S. (1989) 'A beacon navigation method for autonomous vehicles', *IEEE Transactions on Vehicular Technology*, Vol. 38, No. 3, pp.132–139.
- Nasipuri, A. and el Najjar, R. (2006) 'Experimental evaluation of an angle based indoor localization system', *Proceedings of the 4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, 2006.
- Niculescu, D. and Nath, B. (2003) 'Ad hoc positioning system (APS) using AOA', *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pp.1734–1743.
- Priyantha, N.B., Chakraborty, A. and Balakrishnan, H. (2000) 'The cricket location-support system', *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MobiCom)*, pp.32–43.
- Priyantha, N.B., Miu, A.K.L., Balakrishnan, H. and Teller, S. (2001) 'The cricket compass for context-aware mobile applications', *Proceedings of the 7th ACM Conference Mobile Computing and Networking (MOBICOM)*, Rome, Italy.
- Römer, K. (2003) 'The lighthouse location system for smart dust', *Proceedings of the 1st International Conference on Mobile Systems, Applications and Services (MobiSys)*, pp.15–30.
- Rong, P. and Sichitiu, M.L. (2006) 'Angle of arrival localization for wireless sensor networks', *Proceedings of the 3rd Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pp.374–382.
- Roy, R., Paulraj, A. and Kailath, T. (1986) 'Esprit—a subspace rotation approach to estimation of parameters of cisoids in noise', *IEEE Transactions on Acoustics, Speech and Signal Processing*, Vol. 34, No. 5, ISSN 0096-3518.
- Schmidt, R. (1986) 'Multiple emitter location and signal parameter estimation', *IEEE Transactions on Antennas and Propagation*, Vol. 34, No. 3, ISSN 0018-926X.
- Texas Instruments (2007) *CC1000: Single Chip Very Low Power RF Transceiver*. Available online at: <http://focus.ti.com/docs/prod/folders/print/cc1000.html>
- Want, R., Hopper, A., Falcao, V. and Gibbons, J. (1992) 'The active badge location system', *ACM Transactions on Information Systems*, Vol. 10, No. 1, pp.91–102.