

LOCATIONING IN DISTRIBUTED AD-HOC WIRELESS SENSOR NETWORKS

Chris Savarese, Jan M. Rabaey

Berkeley Wireless Research Center
{savarese,jan}@eecs.berkeley.edu

Jan Beutel

Computer Engineering and Networks Lab
ETH Zurich
beutel@tik.ee.ethz.ch

ABSTRACT

Evolving networks of ad-hoc, wireless sensing nodes rely heavily on the ability to establish position information. The algorithms presented herein rely on range measurements between pairs of nodes and the *a priori* coordinates of sparsely located anchor nodes. Clusters of nodes surrounding anchor nodes cooperatively establish confident position estimates through assumptions, checks, and iterative refinements. Once established, these positions are propagated to more distant nodes, allowing the entire network to create an accurate map of itself. Major obstacles include overcoming inaccuracies in range measurements as great as $\pm 50\%$, as well as the development of initial guesses for node locations in clusters with few or no anchor nodes. Solutions to these problems are presented and discussed, using position error as the primary metric. Algorithms are compared according to position error, scalability, and communication and computational requirements. Early simulations yield average position errors of 5% in the presence of both range and initial position inaccuracies.

1. INTRODUCTION

Ad-hoc wireless sensor and actuator networks have many attractive applications. Environmental control and monitoring, smart rooms, robot control, inventory systems, interactive toys, and interactive virtual worlds are only a few examples [1]. Positioning is a key enabler for many of these applications. Sensor data without complete coordinates (this is time stamp, and xyz location) is next to useless. While the Global Positioning System (GPS) offers a solution for localization in an outdoor environment, no such option exists for an indoor setting.

Section 2 of this paper will introduce the general concepts and problems encountered when implementing positioning algorithms in ad-hoc sensing networks. Section 3 introduces the idea of using redundancy to reduce the error introduced by imprecise measurements in a local positioning problem. Section 4 discusses solutions to extend this approach to cover the random topology of a large-scale ad-hoc network.

2. THE POSITIONING PROBLEM IN A MULTIHOP NETWORK

A sensor network typically consists of a large number of nodes with a dense distribution. To reduce the power consumption attributed to communication and to minimize interference, every node can only communicate to its immediate neighbors resulting in a mesh of connections. Clustering as well as the depletion of a local area can occur when mobile nodes move about (see figure 1).

At times it might happen that a partition of a network loses contact to the remaining network due to motion or obstacles blocking the radio signals. To prevent this from happening, the power range of the radios is adaptively set so that each node has a reasonable number of neighboring nodes (first and second order) at any point in time. In other words, it is ensured that the network graph is generally well connected.




Fig. 1. Overlapping ranges of many PicoNodes

The network activity in sensor networks is low and random. Many nodes will be stationary for most of the time, enabling long iteration periods. Within this static framework, a few nodes may move around with limited mobility.

2.1. Navigation

Navigation with radiolocation techniques consists of two components: distance measurements and triangulation. Distance or range measurements can be based on different physical variables: received signal strength (RSSI), angle of arrival (AOA), time of arrival (TOA) or time-distance of arrival (TDOA) of a signal. Three or more independent range measurements with respect to beacon nodes can then be used to solve a 3D-triangulation problem. If these beacons reside at a known location, the absolute position can be given in reference to this inertial system, such as is done for the GPS system [2].

In general, the triangulation problem can be formulated as follows: given a set of references X_i, Y_i, Z_i and a set of range mea-

measurements R_i , a system of linear equations needs to be solved for the unknown U_i .

$$\begin{bmatrix} (X_1 - U_x)^2 + (Y_1 - U_y)^2 + (Z_1 - U_z)^2 \\ (X_2 - U_x)^2 + (Y_2 - U_y)^2 + (Z_2 - U_z)^2 \\ \vdots \\ (X_n - U_x)^2 + (Y_n - U_y)^2 + (Z_n - U_z)^2 \end{bmatrix} = \begin{bmatrix} R_1^2 \\ R_2^2 \\ \vdots \\ R_n^2 \end{bmatrix}$$

Navigation solutions can be classified based on accuracy, availability and cost (hardware, compute cycles, latency and consumed energy). Furthermore, one can differentiate between absolute and relative positioning approaches, where absolute means with respect to a reference point and a map, while relative involves none of those and focuses on topology discovery.

2.2. Localization Challenges in Sensor Networks

When applied to an ad-hoc sensor network, these radiolocation approaches face several new complications: sparse reference points that are not directly visible by all nodes in the network, limited accuracy in the range measurements, and the need for low-power implementation on limited resources. *Anchor nodes*, or nodes with a *prior* knowledge of their locations relative to a global coordinate system, are assumed to be sparse and randomly located. Like the other sensor nodes, their communication range is limited to their immediate neighborhood. This makes it difficult, if not impossible, for *requesting nodes*, or nodes attempting to *estimate* their positions, to acquire enough reference points to perform traditional triangulation. It is only assumed that there will be at least four anchor nodes in a connected network.

The accuracy derived through triangulation depends heavily on the geometry of the position references, the configuration of network nodes, and the accuracy of the range measurements. The short transmit ranges of 1 to 10 m result in unacceptably high synchronization demands of 3 psec per cm of resolution, when TDOA techniques are employed. AOA approaches require costly antenna arrays on each node. These observations make these solutions unattractive, leaving the received signal strength (or RSSI) as the prime candidate for range measurements. Given a known transmission power and a good model of the wireless channel, the distance between transmitter and receiver can be estimated based on the received power. Unfortunately, the accuracy of these RSSI range measurements is highly sensitive to multi-path, fading, non-line of sight (NLOS) scenarios, and other sources of interference, which may result in large errors. These errors can propagate through all subsequent triangulation computations, leading to useless information.

Fortunately, sensor networks possess two properties that may help to overcome these concerns: (i) dense interconnectivity leading to redundancy in the range measurements; (ii) limited mobility which allows for long observation times and the removal of some of the fast-fading effects through integration. In the following section, we first discuss how these properties can be used to solve a local positioning problem (i.e. positioning between nodes that are within communication range). The following section will extend these techniques to a system where not all nodes are within range.

3. LOCAL POSITIONING

3.1. Triangulation

Consider the following scenario: a node with an unknown position receives range measurements (with low accuracy) from a large number (> 3) of neighboring anchor nodes. Using a least-mean squares approach towards solving the over-defined triangulation problem yields a solution with an accuracy that is substantially higher than what could be expected from the unreliable range measurements. Figures 2 and 3 show the simulated positioning results when many ranges are used for a triangulation solution [3].




Fig. 2. Geometric error on triangulations using many nodes

Figure 3 shows a triangulation simulation to 35 independent range measurements that could result from many triangulation updates over time (eliminating fast fading effects) or a dense node population.

3.2. Topology Discovery

It is also worthwhile to consider the inverse problem: a node (with a known position) receives range measurements of a large number of neighboring nodes with unknown position. This information once again can be used to solve a local positioning problem. However, the best that can be accomplished under these conditions is a resolution of the angles between the nodes, or, in other words, the topology of the network. Only relative positioning can be derived. But again, the redundancy in the information helps to increase the resolution of the obtained angles. While a LMS formulation can once again be constructed, we present instead a constructive algorithm that requires only limited computation.

The *Assumption Based Coordinates* (ABC) algorithm determines the locations of unknown nodes one at a time in the order that they establish communication, making assumptions where necessary, and compensating for errors through corrections and redundant calculations as more information becomes available. These assumptions are needed at first in order to deal with the underdetermined set of equations presented by the first few nodes. This description of the general algorithm assumes the perspective of node n_0 .




Fig. 3. Triangulation solutions iterated for 35 independent nodes with independent range errors of 50%

The algorithm begins with the assumption that n_0 is located at $(0, 0, 0)$. The first node to establish communication with n_0 , n_1 , is assumed to be located at $(r_{01}, 0, 0)$, where r_{01} is the RSSI-determined distance between n_0 and n_1 . The location of the next node, n_2 , can then be explicitly solved for, given two assumptions: the square root involved in finding y_2 is assumed to yield a positive result, and z_2 is assumed to be 0.

$$x_2 = \frac{r_{01}^2 + r_{02}^2 + r_{12}^2}{2r_{01}}$$

$$y_2 = \sqrt{r_{02}^2 - x_2^2}$$

The next node, n_3 , is handled much like n_2 , except that only one assumption is made: the square root involved in finding z_3 is positive.

$$x_3 = \frac{r_{01}^2 + r_{03}^2 + r_{13}^2}{2r_{01}}$$

$$y_3 = \frac{r_{03}^2 - r_{23}^2 + x_2^2 + y_2^2 - 2x_2x_3}{2y_2}$$

$$z_3 = \sqrt{r_{03}^2 - x_3^2 - y_3^2}$$

From this point forth, the system of equations used to solve for further nodes is no longer underdetermined, and so the standard algorithm can be employed for each new node. Under ideal conditions, this algorithm thus far will produce a topologically correct map with a random orientation relative to a global coordinate system.

4. GLOBAL POSITIONING

While the above approaches help to improve accuracy in local positioning problems in the presence of unreliable measurements, they do not address the global positioning challenge posed by the ad-hoc wireless sensor networks. Since network-covering beacons represent an unnecessary burden with regards to simple deployment, energy consumption, and network architecture, the only option is to engage in a *cooperative ranging* approach. *Cooperative Ranging* exploits the high connectivity of the network to translate the global positioning challenge into a number of distributed local optimization problems that iteratively converge to a global solution by interacting with each other. The advantage of this approach is that no global resources or communications are needed. The disadvantage is that convergence may take some time and that nodes with a high mobility may be hard to cover. Fortunately, this is not a real issue in sensor networks where nodes rarely move and long discovery times are acceptable given the very long lifetime of the network. This cooperative approach is, to our knowledge, quite original. Existing approaches to localization in sensor networks tend to rely on a global computational engine that receives the range measurements and turns them into an overall optimization problem. An example of such is [4]. The disadvantage of this clever approach are that (i) a global resource is needed challenging the ad-hoc nature of the network, and (ii) that all the range and position information has to be sent back-and-forth to the sensor nodes, resulting in routing bottlenecks and unnecessary energy dissipation.

In the proposed cooperative ranging methodologies, every single node plays the same role, and repeatedly and concurrently executes the following functions:

- Receive ranging and location information from neighboring nodes
- Solve a local localization problem (as introduced in section 3)
- Transmit the obtained results to the neighboring nodes

After the system has converged to a solution, updates are only rarely needed and will be triggered by a mobile node in a localized area of the network.

A number of different cooperative-ranging approaches can be considered. It is worthwhile to differentiate between discovery (startup) and update modes. The former occurs infrequently and last for a short period of time, relative to lifetime of the network, and is responsible for establishing accurate estimates of the stationary nodes in the system. The update mode is invoked after startup, and monitors node movement, updating position information as mobile nodes change their physical location. We will briefly introduce a number of possible approaches and some early results.

4.1. Global Topology Discovery

In this approach, every node assumes initially to be at the center of the coordinate system and performs a local topology discovery (using the ABC algorithm). The resulting information is forwarded to the neighboring nodes. Every anchor node removes a degree of freedom in the coordinate space, and forces the neighboring nodes to linearly transform their own coordinate system (both from a transposition and rotational perspective). This information

is propagated through the system and ultimately causes the system to converge to a single global coordinate space. Early simulations showed that this approach was insufficient to overcome the propagation of initial distance errors, yielding unacceptably large position errors.

4.2. The TERRAIN Approach

The *Triangulation via Extended Range and Redundant Association of Intermediate Nodes* (TERRAIN) algorithm falls in the start-up class. It provides an initial solution for each node in the network by multi-hop forwarding of the anchor positions. At startup, the ABC algorithm is initiated at every anchor node. Requesting nodes wait for the algorithm to propagate to them from at least four independent anchor nodes. A standard triangulation can be performed at that time. As the number of anchor nodes used by each requesting node in this procedure increases, the accuracy of position estimates improves, as expected from the earlier discussion. Note that there is no need to perform the linear transformation at the end of each ABC algorithm to correct for orientation, as a correct topology will provide the needed distance estimate. Figure 4 shows the improvement gained from the TERRAIN algorithm compared to the global topology discovery approach, described above.




Fig. 4. ABC vs. TERRAIN, 32 nodes total, 4 anchor nodes

4.3. Iterative Local Triangulation

Once an initial estimation is obtained, the location accuracy can be improved through an iterative refinement process. Each node uses the most recently computed coordinates of each neighboring node and the range measurements to recompute its own coordinates. This process is iterated several times until the positions of all of the nodes in the network have converged. Figure 5 shows the results of this procedure after 25 iterations in a network cluster of 10 nodes. These results suggest that the accuracy of the final position estimates is influenced more by range errors than by initial position errors. The results shown in figure 4 give an average initial position error of about 39% for TERRAIN at 5% range error. Applying the iterative algorithm afterwards reduces the error to about 5%, an improvement of about 34%.

It should be noted, however, that although the average position errors are low after the refinement stage, simulations show a large




Fig. 5. Refinement, 10 nodes, 4 anchor nodes, 25 iterations

variance in position error values. In a number of cases, the iterative algorithm tends to diverge. The tendency for divergence seems to be correlated with both initial position estimates and range errors; larger sources of error in either area result in higher probabilities of divergence. While rare, further research is necessary to determine techniques to discover and eliminate divergence (for instance by pruning unlikely solutions early on). Reliability of the produced results is an absolute requirement of any sensor-network positioning approach.

5. CONCLUSIONS

Algorithms for positioning nodes in an ad-hoc sensor network have been presented. It has been shown that positioning errors resulting from inaccurate range measurements can be reduced significantly if 7 or more reference points are used in a 3-dimensional triangulation computation [3]. Additionally, simulations show that cooperative ranging, a combination of the TERRAIN and refinement algorithms, is capable of producing position estimates with errors as low as 5%. Further research is required in order to guarantee the stability of this approach.

6. REFERENCES

- [1] Jan M. Rabaey, Josie Ammer, Julio L. da Silva Jr., Danny Patel, and Roundy Shad, "Picoradio supports ad hoc ultra-low power wireless networking," *IEEE Computer*, vol. 33, no. 7, pp. 42–48, July 2000.
- [2] T. Logsdon, *The Navstar Global Positioning System*, Van Nostrand Reinhold, 1992.
- [3] Jan Beutel, "Geolocation in a picoradio environment," M.S. thesis, ETH Zürich, Electronics Lab, 1999.
- [4] Lance Doherty, "Algorithms for position and data recovery in wireless sensor networks," M.S. thesis, University of California at Berkeley, 2000.