Evaluation of Localization Services
Preliminary Report

University of Illinois at Urbana-Champaign
PI: Gul Agha

1 Introduction

As wireless sensor networks (WSNs) scale up, an application's self configurability becomes a critical factor for its successful deployment. In that regard, it is very important to have a resilient yet fairly accurate localization middleware service as many WSN applications rely on the service for self-configuration. Such a localization service aims to determine location information in a given coordinate system for each node in a network being initialized. Several of the methods use *anchor* nodes with known locations and construct relative information about other node positions. Anchor nodes may be hand-placed, have GPS capabilities, or otherwise know their own position in an agreed-upon global coordinate space.

This report presents four sets of localization services with different requirements and different performance characteristics. They are:

1. UIUC service using Vanderbilt's TDoA ranging method.
2. UIUC services using proximity matrices and a maximum likelihood method.
3. Rutgers Ad hoc Positioning System which comes with three versions.
4. Virginia's service employing an approximate point-in-triangle method.

For each set, we describe how localization services work and give a brief evaluation. At the end, we compare the services along key deployment parameters.

2 UIUC Localization Service using TDoA distance measurement

2.1 Description

This localization service from UIUC uses triangulation (also known as trilateration) to determine the coordinate of a node. The implementation assumes a network of anchor and non-anchor nodes. It further assumes that anchor nodes know their location (i.e., coordinate) and are equiped with a sounder; non-anchor nodes have a microphone to receive sound signals from at least three different anchor nodes. At least four anchor nodes should be used to determine the 3-dimensional coordinate of a node. The time difference of arrival (TDoA) service implemented by researchers at Vanderbilt University is used to measure distances needed for triangulation. The method uses the difference in arrival time of radio and sound signals from an anchor node to compute the distance.

Suppose we want to compute the location of a non-anchor node (x, y, z). Let (x_i, y_i, z_i) and d_i be the coordinate of and the distance to the i-th anchor node in range of the non-anchor node, respectively. The method chooses the coordinate (x, y, z) which minimizes the following formula:

$$\sum[(x - x_i)^2 + (y - y_i)^2 + (z - z_i)^2 - d_i^2]^2$$

The gradient descent method is used to find the coordinate minimizing the formula. To compute d_i, the localization service runs the TDoA method several times and takes the median of all the measurements. A CSMA (Carrier Sense Media Access) scheme is used to keep anchor nodes from emitting the sound signal at the same time. That is, before transmitting the sound signal, each anchor node makes sure there is no node sending the sound signal out. Using the median reduces the effect of outliers caused by possible collisions. If a node does not hear from an anchor node for an extended period of time, the anchor node is declared faulty and its corresponding coordinate and distance information are discarded.

2.2 Evaluation

The method gives location estimates with high fidelity even with relatively low node density. Furthermore, the threshold density of anchor nodes to get sufficiently accurate location information is independent of the density of non-anchor nodes in a network. One can improve the accuracy by adding more anchor nodes.

The researchers at Vanderbilt University claim that the ranging method yields 10cm of error given 10m of radio range in an outdoor setting. However, from our indoor experiments with MICA-2 motes, we observed larger errors, even with a shorter radio range. We will contact the Vanderbilt researchers to discuss the causes of the discrepancy and resolve this issue.

Another observation is that the accuracy of distance measurement seems to depend heavily on the orientation of the sounder at an anchor node. Our experiments showed estimated distances vary from .5m to 2m depending on the orientation. This orientation dependence has a serious negative effect on indoor measurements. For example, when a non-anchor node is placed behind a sounder (i.e., speaker), it may hear only echoes and compute an inordinately large distance.

Vanderbilt's ranging service implementation currently suspends the radio While a node emits or hears the beacon sound. The duration of the suspension lasts as long as 1.6 sec. This radio suspension may significantly decrease the available network bandwidth when the localization service is turned on.

3 UIUC Proximity Matrix and Maximum Likelihood Methods

3.1 Description

Both the proximity matrix method and maximum likelihood method are based on a node's neighborhood information. A node's neighbors can send a packet to the node in a single hop. As such, two methods are based only on a node's radio range; they do not use any other sensors or devices.

The proximity matrix method [2] uses proximity matrices; a node's proximity matrix encodes the relative closeness among its neighbor nodes including itself. The idea of the proximity matrix method is based on the observation that a neighbor set from a neighboring node partitions the

neighbor set of the recipient into the intersection and difference sets. If two nodes are close together, the chance that they belong to the same partition is high. This observation lets us use the number of times that two nodes fall in a partition to represent their relative closeness among a node's neighbors. A node computes its proximity matrix by exchanging its neighbor information with its neighbors.

Once the proximity matrix is obtained, a node constructs its local coordinate system by arbitrarily choosing coordinate axes and assigns local coordinates to its neighbors accordingly. Then the coordinate system is rotated to make it consistent with that of the node's parents. As the last step, a polar coordinate system is derived which has the center node as its origin and uses the hop count from the center node as a node's radius; all nodes are assigned global coordinates accordingly.

The maximum likelihood method is similar to the proximity matrix method in that it first computes a local coordinate system for each node and then derives a global coordinate system using the local ones starting from the center node. For each node, the method assigns a local coordinate to each neighboring node in such a way that the assignments maximize the likelihood of the relative node positioning described in the neighbor information exchanged among the nodes.

Once all the nodes compute their local coordinates, the global coordinate system is obtained as follows. First, the center node broadcasts the coordinate assignments to its neighbors. The neighbors align their local coordinate systems with that of the center node by using the coordinates it sent; the process is repeated until all the nodes are consistently assigned a global coordinate.

The two methods compute locations in a relative coordinate system: a central node is assigned or elected to the origin of the coordinate system and it chooses the axes arbitrarily. However, the coordinate system that is obtained can be transformed to an absolute one with additional information.

3.2 Evaluation

The advantages of the maximum likelihood estimation and the proximity matrix methods are that they do not assume the availability of anchor nodes and they do not require other sensors. This makes them usable in a range of deployments.

A drawback of the methods is that they require high node density to get reasonably good accuracy. Another problem is that our current implementations do not precisely handle global coordinate assignment, although they work very well for relative coordinates. However, since these methods estimate the distances within a single message hop, they can also be used to further refine location estimations based on hop counting. Some preliminary experiments have been done with the proximity matrix on a small-scale (10-11) network of MICA-1 motes. These results suggest that the error in estimated position may be in the range of 10cm to 50cm in an indoor environment. The likelihood estimation method has been tested only on a simulator. We are currently porting the two methods to the MICA-2 platform.

4 Localization Services from the Rutgers University

4.1 Description

The Ad-hoc Positioning System from Rutgers University is a family of distributed localization algorithms based on anchor nodes. All non-anchor nodes can be placed at arbitrary locations;

however, they must be able to contact (through one or more hops) at least three different non-colinear anchor nodes. Given distances to and absolute positions of such anchors, all nodes can calculate their own absolute position via triangulation.

There are two major choices for describing relative node location: the minimum number of network hops between nodes and direct physical distance. Researchers at Rutgers University have developed four different methods for propagating this proximity information in the APS framework, called DV-hop, DV-distance, Euclidean and DV-coordinate. The DV-hop method determines the number of network hops from each node to three nearby anchors (i.e. a distance vector), and then scales this metric based on the ratio of hop counts to true distances between anchor nodes; DV-distance uses the sum of measured neighbor-to-neighbor distances in place of hop count. The Euclidean method uses distance measurements and triangle identities to calculate the distance between a node A and an out-of-range anchor L via a pair of nodes B and C which already know their distances from each other and both of A and L. Finally, the DV-coordinate method builds local coordinate systems around each node based on local distance measurements, and then propagates anchor locations, transforming them from one local coordinate system to the next.

4.2 Evaluation

Currently the APS service for TinyOS 1.x implements only the DV-hop propagation method; this places it at a disadvantage in terms of practical deployment for NEST applications. Although Rutgers has not tested their implementation on the Mica-2 platform, the code does not appear to depend on any features which differ between Mica-1 and 2. According to output from NesC, APS together with TinyOS uses 485 bytes of RAM on a mote. We estimate that APS's anchor position and neighbor distance tables, message buffer, and other data consume around 150 bytes of RAM. Although this is fairly tight for the Mica-1's 512 bytes of RAM, it is not as much of an issue on the Mica-2, especially considering that localization only needs to be run at startup, and possibly occasionally thereafter. We were able to run the service on a small number of motes. However, the simulations in the designers' own publications may give a better idea of the large-scale behavior.

Rutgers researchers have compared the four different variants of APS in simulation. In one study [4], 200 nodes are placed with average degree 9; in another [3], 100 nodes are placed with average degree 7.6. Furthermore, both studies examine isotropic and anisotropic node distributions; the former case resembles an open field, while the latter is more representative of the interior of a building, where barriers may block communication between nearby nodes. In all cases position error was measured in terms of radio range.

Although these node degrees may be on the high end for NEST demonstrations, they are not unreasonably large. In general, decreasing node degree will increase the deviation between the direct path as the crow flies and the hop-by-hop path a message follows from an anchor to another node. As a result, error will grow when using the DV-hop and DV-distance techniques. This is similar to the error encountered in the anisotropic case. In [4], DV-hop achieved positioning with an error of 30-40% of radio range in the isotropic case and error of 100-150% in the anisotropic case. DV-distance attained somewhat better position accuracy (isotropic: 10-20%, anisotropic: 90-130%) for low range measurement error (0-25%); as measurements get worse, DV-distance becomes similar to DV-hop in the anisotropic case and worse in the isotropic case, with position errors 50-75% of radio range.

The Euclidean and DV-coordinate approaches do not show significantly different performance between the isotropic and anisotropic cases. This stands to reason because they explicitly take into

account direction via second-hop information. For small numbers of anchors (less then 10% of all nodes), the Euclidean method's error exceeds DV-hop's in the isotropic case. However, increasing the number of anchors to 20-30%, error decreases to 10-20% of radio range given low ranging error (0-25%), or at worst about 50%. This is similar to the performance of DV-distance in the isotropic case, except the Euclidean method is also able to maintain this performance in the anisotropic case. The DV-coordinate approach demonstrates very high position error even for low ranging errors, e.g. position error 75-125% of radio range for ranging error less than 10%. In some cases, position error exceeds twice the radio range. This is due to the rapid accumulation of errors in the anchor positions transmitted through the network as they are transformed from one local coordinate system to another. Furthermore, the error is slightly worse in the anisotropic than the isotropic case because the average number of hops between nodes (i.e. opportunities to introduce error) is greater.

Based on the above discussion, we believe the APS service currently available for TinyOS will not be applicaple to the indoor NEST demo scenarios. The anisotropic environment is at odds with the DV-hop method's assumption of good correlation between hop count and distance. This service may, however, be reasonable for an outdoor (or single large room) environment, provided that nodes have sufficiently many neighbors to determine a reasonably straight path to the anchors. An interesting possibility is to use the Rutgers DV-Hop service to reinforce the results of another localization service which does use ranging measurements.

5 Localization Service from the University of Virginia

5.1 Description

The Approximate Point-In-Triangulation (APIT) localization algorithm [1] is an area-based, range-free localization method. Given that a small fraction of the nodes in the sensor network are location-aware *anchors*, this algorithm computes locations for other nodes in the network by determining the smallest possible area in which the node may be located. This can be accomplished via the point-in-triangulation (PIT) test – a geometric technique for testing whether a point is inside the given triangle. Since the PIT test needs the distances between points, it cannot be used directly in a range-free algorithm. Instead the approximate PIT test is applied, which relies only on the monotonicity of the reduction in radio signal strength along a single direction. APIT allows a node to determine whether it is closer or farther than another node to the source of the signal.

Localization with the APIT test is performed as follows: a point performs APIT tests for each possible triangle whose vertices are anchors heard by the node. The intersection of the triangles satisfying the APIT test is the approximate location of the node, and the center of gravity of that area is taken as the estimate of the location.

5.2 Evaluation

The principal advantage of this approach is that it does not rely on any sensor- or radio-based ranging techniques, which can be inaccurate and are affected by environmental conditions. Only the most basic properties of radio wave propagation are assumed. The main disadvantage is the need for a large number of anchor nodes within the vicinity of each node, which is necessary to perform multiple APIT tests. Because of this requirement, the algorithm's precision depends heavily on the node density and radio range.

The algorithm has only been tested in simulation. There is no TinyOS implementation.

6 Comparison

Operating Environment

	UIUC/Vanderbilt	UIUC PM/ML	Rutgers	Virginia
# of nodes	unlimited	unlimited	unlimited	unlimited
# of anchors	≥ 3, depends on the acoustic sensing range	None	5-33% anchors for DV-hop/DV-distance, 15-33% anchors for Euclidean	10-20 in radio range of each node
node density (neighbors per node)	depends on the acoustic sending range	10-15	7-10	> 10
radio range		uniform	small/large	large
obstacles	depends on echoes/obstacles	unaffected	bad (DV-hop/-distance) good (Euclidean)	unaffected
node distribution	unaffected	uniform distribution	nodes being along a straight line from anchors to non-anchors	unaffected

Precision

UIUC/Vanderbilt	UIUC PM/ML	Rutgers	Virginia
>10cm error in a 10m range (outdoors) <2m indoor	N/A	10-50cm on MICA-1 (under good condition)	depends on radio range/ # of anchors heard (can achieve < .2R)

API Compliance

UIUC/Vanderbilt	UIUC PM/ML	Rutgers	Virginia
Yes. (StdControl, getLocation, locationReady; setLocation for anchors)	N/A (simulation only)	Yes, in latest release.	N/A (simulation only)

Memory Footprint

UIUC/Vanderbilt	UIUC PM/ML	Rutgers	Virginia
18976B ROM 2557B RAM	N/A	458B RAM (APS/DV-hop) 150B RAM (its data alone)	N/A

Sensors Used

UIUC/Vanderbilt	UIUC PM/ML	Rutgers	Virginia
speaker, microphone, and radio	Radio only	Radio only	Radio only

Environment

UIUC/Vanderbilt	UIUC PM/ML	Rutgers	Virginia
Outdoor (currently working on improvment in indoor settings)	PM assumes nodes in a 2-D plane	Indoor & outdoor (for currrent implementation: no obstruction)	Indoor & outdoor

References

[1] Tian He, Chengdu Huang, Brian M. Blum, John A. Stankovic, and Tarek Abdelzaher. Range-free localization schemes for large scale sensor networks. Technical Report CS-2003-06, University of Virginia, Computer Science Department, March 2003.

[2] YoungMin Kwon, Wooyoung Kim, and Gul Agha. Self localization of wireless smart micro sensors using proximity matrices. (submitted for publication), 2003.

[3] Dragos Niculescu and Badri Nath. Ad-hoc positioning system (APS). In *Proceedings of Symposium on Ad-Hoc Wireless Networks, GlobeCom2001*, San Antonio, Texas, November 2001. also Technical Report DCS-TR-435, Rutgers University, April 2001.

[4] Dragos Niculescu and Badri Nath. DV-based Positioning in Ad hoc Networks. *Telecommunication Systems*, 22(1-4):267–280, 2003.