Ad Hoc Positioning System (APS) Using AOA

Dragoş Niculescu and Badri Nath
DATAMAN Lab
Rutgers University
email: {dnicules,badri}@cs.rutgers.edu

Abstract— **Position information of individual nodes is useful in implementing functions such as routing and querying in ad-hoc networks. Deriving position information by using the capability of the nodes to measure time of arrival (TOA), time difference of arrival (TDOA), angle of arrival (AOA) and signal strength have been used to localize nodes relative to a frame of reference. The nodes in an ad-hoc network can have multiple capabilities and exploiting one or more of the capabilities can improve the quality of positioning. In this paper, we show how AOA capability of the nodes can be used to derive position information. We propose a method for all nodes to determine their orientation and position in an ad-hoc network where only a fraction of the nodes have positioning capabilities, under the assumption that each node has the AOA capability.**

Index Terms— **ad hoc networks, positioning, orientation, digital compass, AOA.**

I. Introduction

The main features of new adhoc networks are large number of unattended nodes with varying capability, lack or impracticality of deploying supporting infrastructure, and high cost of human supervised maintenance. What is necessary for these types of networks is a class of algorithms which are scalable, tunable, distributed, and easy to deploy. With recent advances in small device architectures [1], it can be foreseen that cheap, or even disposable nodes, will be available in the future, enabling an array of new agricultural, meteorological and military applications. These large networks of low power nodes face a number of challenges: cost of deployment, capability and complexity of nodes, routing without the use of large conventional routing tables, adaptability in front of intermittent functioning regime, network partitioning and survivability. It is a given that in many of these networks, due to considerations of cost, size, and power requirements, individual nodes will have varying capabilities. A general question is how to export capabilities to various nodes in the network so that the overall capability can be increased in the network. For example, many ad hoc network applications and protocols assume the knowledge of geographic position of nodes. The absolute position of each networked node is an assumed fact by most sensor networks which can then present the sensed information on a geographical map. Also, the availability of position would enable routing in sufficiently isotropic large networks, without the use of large routing tables.

However, not all nodes have the capability of locally determining their position by means of GPS. Finding position without the aid of GPS in each node of an ad hoc network is important in cases where the GPS service is either not accessible, or not practical to use due to power, form factor or line of sight conditions such as indoor sensors, sensors hidden under foliage, etc. A similar argument holds for orientation as compasses face erratical behavior in the vicinity of large metal objects or electrical fields. Orientation, or heading, is used in remote navigation, or remote control of specialized sensors, such as directional microphones or cameras. In this paper, we address the problem of self positioning and orientation of the nodes in the field, which may provide a general framework for exporting capabilities in a network where more capable nodes cooperate in dispersing information to less capable nodes.

What is necessary for ad hoc deployment of temporary networks is a method similar in capability to GPS and magnetic compasses, without requiring extra infrastructure, or extensive processing capabilities. What we propose is a method by which nodes in an ad hoc networks collaborate in finding their position and orientation under the assumptions that a small fraction of the network has only the position capability. A compass is not necessary in any node, but if it is available, either at the landmarks, or everywhere, it will enhance the accuracy of the positioning algorithm. Previous positioning methods used so far used either TDOA, like in Cricket[2] and AhLOS[3], or signal strength (RADAR[4], APS[5]). What makes our approach different from previous ones is that it is based on the capability of the nodes to sense the direction from which a signal is received, which is known as angle of arrival (AOA). AOA sensing requires either an antenna array, or several ultrasound receivers, but besides positioning, it also provides the orientation capability. This is currently available in small formats in wireless networked nodes such as the one developed by the Cricket Compass project[6] from MIT. In fact, APS using AOA is part of a larger effort to provide positioning based on multimodal sensing. The

This research work was supported in part by DARPA under contract number N-666001-00-1-8953

aim is to show that ad hoc positioning is possible using various localization capabilities (ranging, AOA, compasses), independently, or together.

One scenario involving sensor networks frequently mentioned in literature is that of aircraft deployment of sensors followed by in flight collection of data by simply cruising the sensor field. This and other meteorological applications are implicitly assuming that the data provided by the sensor is accompanied by the sensor's position. It is thus possible to attach the sensed information to a geographical map of the monitored region. If this is an absolute necessity for making sense of the observed data, accurate position might also be useful for routing and coordination purposes. For some ad hoc networks, algorithms such as Cartesian routing[7], or geocast[8], enable routing with reduced or no routing tables at all, and are appropriate for devices like the Rene mote[1], with only half a kilobyte of RAM. An improvement that can be applied to some ad hoc routing schemes when position is available, Location Aided Routing [9] limits the search for a new route to a smaller destination zone. Our positioning and orientation algorithm is appropriate for indoor location aware applications, when the network's main feature is not the unpredictable, highly mobile topology, but rather temporary and ad hoc deployment. These networks would not justify the cost of setting up an infrastructure to support positioning, like proposed in [10], [4], or [2].

The orientation and positioning problems have been extensively studied in the context of mobile robot navigation [11]. However, many methods proposed by the robotics community make extensive use of image processing and preset infrastructure, such as "recognizable" landmarks. Our aim is a positioning method that is robust, but relies on less computational resources and infrastructures.

The rest of the paper is organized as follows: the next section describes the assumptions of the problem and the basic properties of AOA capable nodes. Section III presents our proposed approach, with the orientation forwarding scheme, and section IV discusses some error control issues. Sections V and VI present simulation results and discuss some mobility related issues, and VII summarizes with some concluding remarks.

II. AOA THEORY

The network is a collection of ad hoc deployed nodes such that any node can only communicate directly with its immediate neighboring nodes within radio range. In the ideal case, when radio coverage of a node is circular, these networks are modeled as fixed radius random graphs. Each node in our network is assumed to have one main axis against which all angles are reported and the capacity to estimate with a given precision the direction from which a neighbor is sending data. We assume that after the deployment, the axis of the node has an arbitrary,

Fig. 1. Nodes with AOA capability

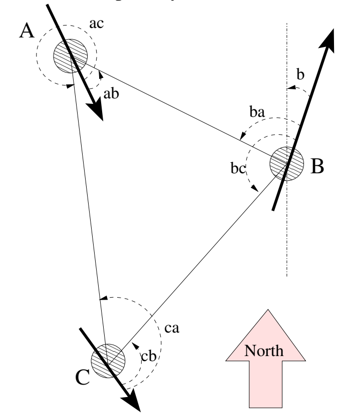

unknown heading, represented in figure 1 by a thick black arrow. Some of the nodes, from here on called landmarks, have additional knowledge about their position from some external source, such as a GPS receiver or human input. The term *bearing* refers to an angle measurement with respect to another object. In our case, the AOA capability provides for each node bearings to neighboring nodes with respect to a node's own axis. A *radial* is a reverse bearing, or the angle under which an object is seen from another point. We will use the term *heading* with the meaning of bearing to north, that is, the absolute orientation of the main axis of each node. In figure 1, for node B, bearing to A is $\widehat{ba}$, radial from A is $\widehat{ab}$, and heading is $\widehat{b}$. The problem to be solved is: *given imprecise bearing measurements to neighbors in a connected ad hoc network where a small fraction of the nodes have self positioning capability, find headings and positions for all nodes in the network*. The difficulty of the problem stems from the fact that the capable nodes (landmarks) comprise only a small fraction of the network, and most regular nodes are nodes are not in direct contact with enough landmarks. What we are looking for is a hop by hop method to export capabilities from the capable nodes to the regular ones.

When interacting with two neighbors, as shown in figure 1, a node can find out the angle between its own axis and the direction the signal comes from. Node A "sees" its neighbors at angles $\widehat{ac}$ and $\widehat{ab}$, and has the possibility of inferring one angle of the triangle, $\widehat{CAB} = \widehat{ac}$-$\widehat{ab}$. For consistency all angles are assumed to be measured in trigonometric direction. Node A can also infer its heading, if heading of one of the neighbors, say B, is known. If node B knows its heading (angle to the north) to be $\widehat{b}$, then A may infer its heading to be $2\pi - (\widehat{ba} + \pi - \widehat{ab}) + \widehat{b}$. This in fact a way to export the compass capability from B to A. If however, no compass is available in any node, but each node knows its position, heading can still be found because the orientations for the sides of the triangle can be found from positions of its vertices. What is therefore needed is

Fig. 2. Basic principle of Cricket compass (adapted from[6]).

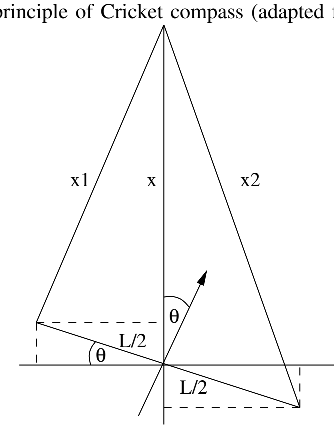

a positioning algorithm based on AOA, orientation being available by one of the means mentioned above.

A. AOA capable nodes

AOA capability is usually achieved by using an antenna array, which might be prohibitive in size and power consumption. A small form factor node that satisfies conditions outlined in the previous section has been developed at MIT by the Cricket Compass project[6]. Its theory of operation is based on both time difference of arrival (TDoA) and phase difference of arrival. Time difference is used in a similar manner in other projects, such as AhLOS[3] and Cricket Location[2], and is based on the six orders of magnitude difference between the speeds of sound and light. If a node sends a RF signal and an ultrasound signal at the same time, the destination node might infer the range to the originating node based on the time difference in arrivals. In order to get the angle of arrival, each node may use two ultrasound receivers placed at a known distance from each other, L (figure 2). By knowing ranges x_1, x_2, and distance L, the node is able to infer the orientation θ, with a accuracy of $5°$ when the angle lies between $\pm 40°$. Medusa, used in AhLOS project[3] from UCLA, is another wireless networked node with small size which makes use of several ultrasound receivers, but without actually employing them to detect angle of arrival. These incipient realizations prove that it is feasible to get AOA capability in a small package that would be appropriate for future pervasive ad hoc networks.

B. Triangulation Using AOA

The central observation suggesting that positioning using AOA is possible is that the following: if we know the positions for the vertices of a triangle and the angles at which an interior point "sees" the vertices, we can determine the position of the interior point. This problem, called triangulation, is somewhat similar to the trilateration problem, used in GPS[12]. The difference is that the interior point knows angles towards triangle sides instead

Fig. 3. Positioning by measuring angles to landmarks

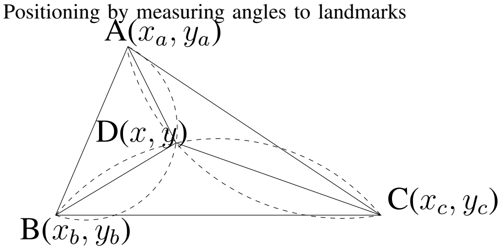

of distances to vertices. In figure 3, if beside coordinates of A, B and C, node D knows distances DA, DB and DC, it can use trilateration to infer its position.On the other hand, if it knows the angles $\widehat{BDA}$,$\widehat{ADC}$, and $\widehat{CDB}$ it can find its position using triangulation. This is done by finding the intersection of the three circles determined by the landmarks and the known angles. Information from several landmarks can be used to get a least square error solution, because in the general case, AOA measurements do not have perfect accuracy. There are several possibilities to compute this estimated point of intersection. For this explanation it is useful to review how positioning is done using trilateration, when distances to known landmarks are known. Given (x_i, y_i) and d_i, the coordinates of and respectively the distance to landmark i, we build the nonlinear system

$$\begin{aligned}(x - x_i)^2 + (y - y_i)^2 = d_i^2 \\ i = 1, ..., n\end{aligned}$$

In Global Positioning System[12], the system is solved using nonlinear methods based on successive approximations, but it also can be solved by reduction to a linear system by subtracting one equation from the rest. In this latter case, we obtain equations

$$\begin{aligned}&2x(x_i - x_1) + 2y(y_i - y_1) = \\ &= d_1^2 - d_i^2 + x_i^2 - x_1^2 + y_i^2 - y_1^2 \quad , \ i = 2, ..., n\end{aligned}$$

This linear system can be solved using standard methods for over-determined systems, such as the pseudo-inverse.

Getting back to our triangulation problem, it can be reduced to trilateration by some simple transformations. If for example a node D knows the angle to a pair of landmarks A and B, it may infer that its position is somewhere on the circle determined by the angle and the position of the two landmarks (figure 4). What is fixed in this picture is the center of the circle, O, whose position may be determined when x_a, y_a, x_b, y_b and angle $\widehat{ADB}$ are known. This may help in transforming a triangulation problem of size n into a trilateration problem of size $\binom{n}{2}$ if for each pair of landmarks observed by a node we create an trilateration equation using x, y, x_0, y_0 and the radius of the circle as the distance. Another possibility is to form all triplets of obtained landmarks and find the center of the

Fig. 4. Geometric locus of a point seeing two known landmarks at a given angle, is a circle. $\widehat{AOB} = 2\widehat{ADB}$

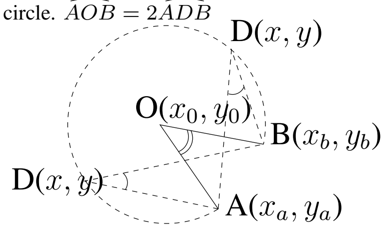

circumscribed circle for each such triplet and the unknown point D. This leads to the solving of $\binom{n}{3}$ trilateration problems of size 3, one for each circle, with $\binom{n}{3}$ solution points. For small numbers of landmarks, the $\binom{n}{3}$ has a similar CPU complexity as the $\binom{n}{2}$ one. However, the problems solved are all of size 3, thus requiring much less memory, whereas the $\binom{n}{2}$ approach needs to handle $\frac{n^2}{2}$x2 sized matrices. A solution linear in the number of landmarks n, proposed in [11], makes efficient use of the representation of landmarks as complex numbers. In our simulation we used the simple $\binom{n}{3}$ implementation, as it gives the same quality estimates as the linear solution presented in [11], but it is much more simple to implement and has a low penalty for small n.

Another method of positioning using angles is VOR (VHS Omni-directional Range), which is currently still the main aid for aircraft navigation. Its principle is very simple: a landmark sends two signals, one that is periodic and omni-directional, while the another one is directional and is rotated about the landmark. The airborne equipment receives both signals, and interprets the difference as a radial from the station. The coordinates of the station are known, therefore placing the mobile anywhere on a given line. A second VOR reading provides a second line to be intersected with the first. Given (x_i, y_i, r_i) the coordinates and the radial to the landmark i, a node can build the equation of the line $a_i x + b_i y = c_i$ on which it places itself.

$if\ \cos(r_i) = 0$

$\quad a_i = 1;\ b_i = 0;\ c_i = x_i$

$else$

$\quad a_i = \tan(r_i);\ b_i = -1;\ c_i = -y_i + x_i \tan(r_i)$

Combining all such lines to landmarks, the linear system to be solved for a location is:

$$\left[\begin{array}{c|c} a^T & b^T \end{array}\right] \begin{bmatrix} x \\ y \end{bmatrix} = \left[c^T\right]$$

This approach is less expensive computationally, for n landmarks requiring just a weighted least square linear system solving. What makes it slightly different form the previous one is the fact that the landmark should be equipped with a compass, so that it reports all radials against a well known direction, such as north. The bearing method, on the other hand does not require any compass at all, but still provides positioning and orientation for the

Fig. 5. Node A infers its bearing to L by corroborating $B's$ and $C's$ bearings to L

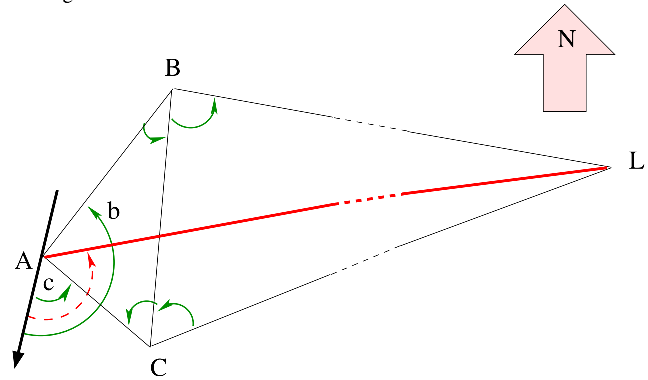

all nodes.

III. Ad Hoc Positioning System (APS) algorithm

The problem in an ad hoc network is that a node can only communicate with its immediate neighbors, which may not always be landmarks (landmarks are nodes which know their position and possibly their heading). APS [5] is a hybrid between two major concepts: distance vector (DV) routing and beacon based positioning (GPS). What makes it similar to DV routing is the fact that information is forwarded in a hop by hop fashion, independently with respect to each landmark. What makes it similar to GPS is that eventually each node estimates its own position, based on the landmark readings it gets. The original APS concept has been shown to work using range measurements, but is in fact extensible to angle measurements. What we propose is a method to forward orientation so that nodes which are not in direct contact with the landmarks can still infer their orientation with respect to the landmark. Here, "orientation" means bearing, radial, or both. We examine two algorithms, *DV-Bearing,* which allows each node to get a bearing to a landmark, and *DV-Radial,* which allows a node to get a bearing *and* a radial to a landmark.

The propagation works very much like a mathematical induction proof. The fixed point: nodes immediately adjacent to a landmark get their bearings/radials directly from the landmark. The induction step: assuming that a node has some neighbors with orientation for a landmark, it will be able to compute its own orientation with respect to that landmark, and forward it further into the network. What remains to be found is a method to compute this induction step, both for bearings and radials.

A. Orientation Forwarding

The method is shown in figure 5: assume node A knows its bearings to immediate neighbors B and C (angles $\widehat{b}$ and $\widehat{c}$), which in turn know their bearings to a faraway landmark L. The problem is for A to find its bearing to L

(dashed arrow). If B and C are neighbors of each other, then A has the possibility to find all the angles in triangles ΔABC and ΔBCL. But this would allow A to find the angle $\widehat{LAC}$, which yields the bearing of A with respect to L, as $\widehat{c}+\widehat{LAC}$. Node A might accept another bearing to L from another pair of neighbors, if it involves less hops than the pair $B-C$. A then continues the process by forwarding its estimated bearing to L to its neighbors which will help farther away nodes get their estimates for L. Forwarding orientations is done in a fashion similar to distance vector routing algorithms. In our case, the landmarks are the ones starting the update messages that are propagated throughout the network, for each landmark independently. Once node A finds its bearings to at least three landmarks that are not on the same line or on the same circle with A, it can infer its position using one of the methods outlined in section II-B.

If the radial method is to be used, a similar argument holds, with the difference that now A needs to know, besides bearings of B and C to L, the radials of B and C from L. If the angle $\widehat{BLN}$ (radial at B) is also known, then the angle $\widehat{ALN}$ (radial at A) can also be found since all angles in both triangles are known. The actual downside for this method is in the increased amount of signaling - nodes B and C forward two values per landmark (bearing and radial) instead of just one, as in the bearing based method. If a compass would be available in every node, the two methods would in fact become identical because when all angles are measured against the same reference direction (north, for example), $bearing = \pi - radial$.

The algorithm has similar signalling overhead behavior with the original APS [5] algorithm (range based), which is roughly a TTL limited flooding per landmark. The following table summarizes for each method the required node capabilities and associated signaling-accuracy trade-offs. "More" signalling refers to the fact that two values are needed per landmark, whereas "less" sends only one. In the case of a large existing packet overhead, one extra value may be of diminished importance. The accuracy of the two propagated methods will be quantified more precisely in the simulation section (V).

compass	method	signaling	accuracy
nowhere	*DV-Bearing*	less	less
only at landmarks	*DV-Bearing*	less	less
	DV-Radial	more	more
all nodes	*DV-Radial*	less	more

B. Network density

The question that arises in deployment of the network is what kind of node density is needed in order to achieve a certain condition with high probability. It has been conjectured[13] that many random graph properties exhibit phase transitions - sharp increases in the probability when the density increases beyond a certain point. For example,

Fig. 6. Probability for a node to satisfy conditions necessary for orientation forwarding

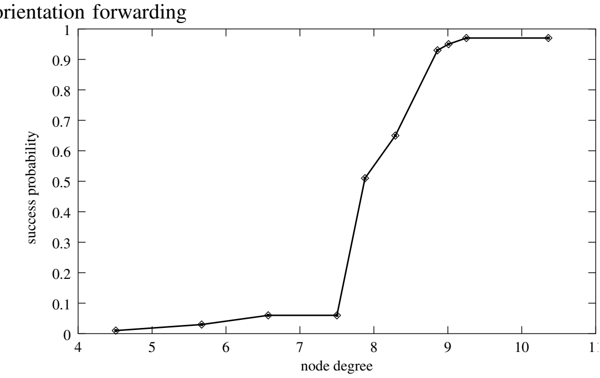

it has been proven that a random network in the plane needs a degree of about 6 in order to have complete connectivity with high probability. We expect the degree requirement to be higher in our case, since more than simple connectivity is needed for the orientation propagation to work. For a bearing to be propagated, two neighbors that are also neighbors of each other should be present for any given node. In figure 6, we can see that when the mean degree of a node increases beyond 9, with a very high probability it will meet the conditions to forward orientation. This data is empirically obtained by running our forwarding policy in a network of 1000 nodes with a single landmark and then count the number of nodes which get a bearing to it. Variation of the average degree is achieved by increasing the radio range of the nodes. In the case of a sensor network, it is often envisioned that the deployed density is higher than needed to allow for extension in battery life by tuning the duty cycle. This means that an initial degree of 9 might be tolerable (50% more nodes have to be deployed), as the normal functioning regime can be later lowered to 6, which has been shown to be the minimum for connectivity [13].

IV. Error control

Being that all bearing measurements are affected by errors, the forwarding may actually amplify and compound smaller errors into larger errors. A number of simple techniques may be employed to reduce the propagation of such errors, including: avoiding inference based on small angles or on degenerate triangles, limiting the propagation of DV packets with a simple TTL scheme and elimination of the outliers in position estimation. The fact that angle measurements are affected by error greatly influences the very core of our algorithm: bearing propagation. As the environment we envision for this positioning algorithm is a low power, low communication capacity node, error control methods employed have to be lightweight. Together, the three mentioned methods achieve an error reduction of about half.

Fig. 7. Error is compounded with propagation

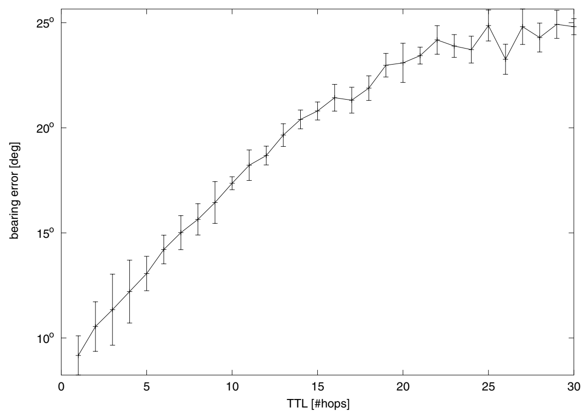

Fig. 9. Removing the outliers from centroid computation

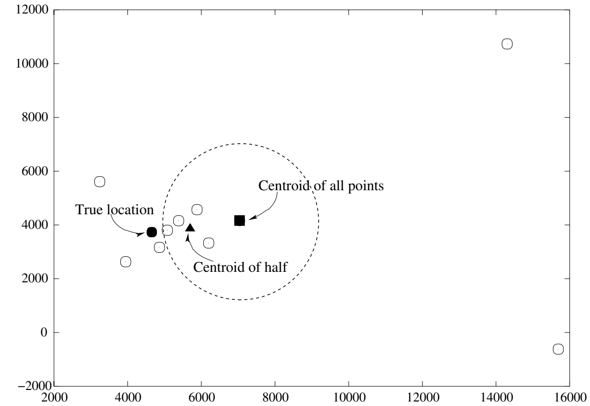

Fig. 8. Angle threshold provides a tradeoff between error and coverage

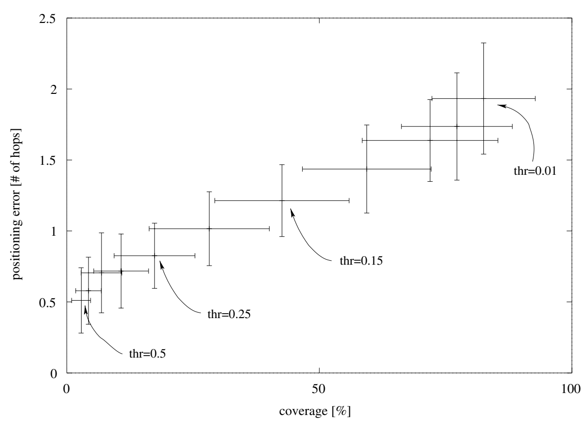

The first intuitive remark is that error cumulates with distance, because of the way bearings are propagated. We verified this fact in a network of 200 nodes, by plotting the average bearing error as a function of distance in hops to the respective landmark (figure 7). Limiting the propagation of the DV packets using a TTL scheme is a good idea not only for error control reasons, but also for reducing communication complexity. If TTL is infinite, each landmark is flooding the entire network with its coordinates, thus triggering bearing computation at every other node. Therefore, the TTL is the main feature that makes the proposed algorithm scalable. As long as enough landmarks can be acquired from the area allowed by the TTL, the total size of the network does not influence the amount communication or the quality of the estimates.

The next key observation is that small angles are more error prone than large angles. It is preferable to deal with equilateral triangles than with triangles that have two very acute angles and one obtuse angle. The reason for that is that AOA measurement errors are theoretically independent of the actual angle measured. For example, the same angle error of 5^o will make much more difference in a triangle with a true angle of 10^o than in one with 60^o. It is an analogue situation with the geometric dilution of precision (GDOP) in the GPS, in which the error amplification depends on the landmark constellation. To address this problem we use a threshold value to eliminate triangles in which small angles are involved. In figure 5, section III-A, if any of the two triangles ΔABC and ΔBCL have small angles, A won't get to propagate orientation to L. There is a tradeoff between coverage and positioning error, and this results from the orientation forwarding policy. A conservative policy would use a high threshold, limiting the computations with small angles but also limiting the propagation of orientations, and finally reducing coverage. A relaxed policy would propagate almost all angles, involving more errors, but would improve coverage. In figure 8, it is shown how varying from a very conservative forwarding policy (threshold=0.5$\simeq 28°$) to a very relaxed one (threshold=0.01$\simeq 0.5°$) achieves different levels of coverage (success rate) with different amounts of error. The positioning error represents the average distance in hops from the true position, obtained after propagation (section III-A) and triangulation (section II-B). The coverage represents the fraction of nodes successfully obtaining a position. The data is obtained by positioning with different thresholds, indicated in the figure, in the configuration shown in figure 10, with a TTL of 6, 20% landmarks and a high AOA measurement error (stddev = 0.4 $\simeq 23°$).

Another error control method, suggested in [11], refers to the position estimations obtained from the triplets of landmarks. In all the mentioned methods, several position estimations may be obtained, leading to the problem of combining them into one single estimate. While this can simply be the centroid of the estimates, in practice it has been observed that large errors are clustered together. This is caused by common angle errors across bearing propagation paths. The method suggested in [11] is to first compute the centroid and then remove the outliers before

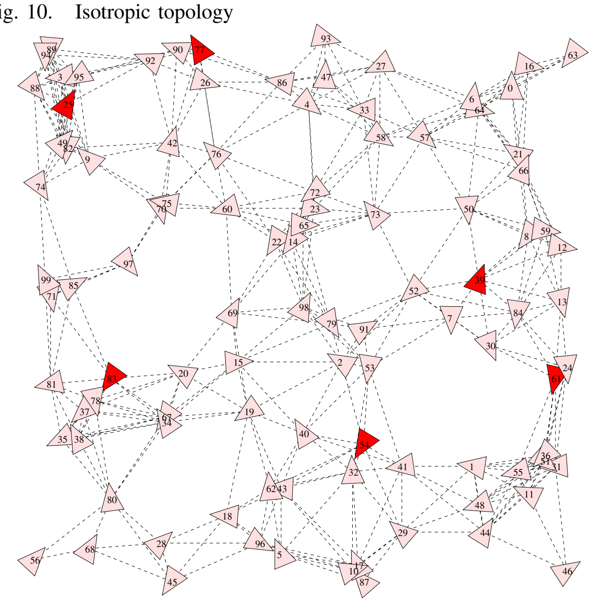

Fig. 10. Isotropic topology

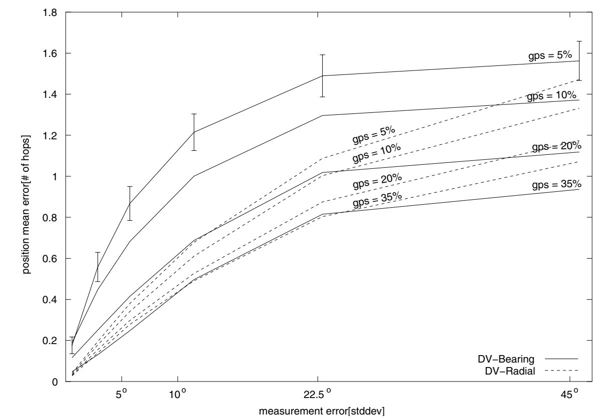

Fig. 11. Positioning error

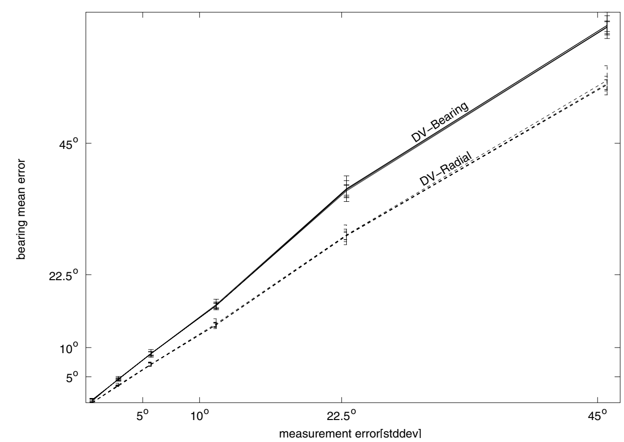

Fig. 12. Bearing error

recomputing a new centroid with the remaining points (fig 9). There are more powerful methods available, such as data clustering and k-smallest enclosing circle, but they involve higher computational and memory complexities, which may not be applicable to small networked nodes, such as sensor nodes.

V. SIMULATION

We simulated an isotropic[1] map similar with the one in figure 10 (average degree=10.5, diameter=32), but with 1000 nodes, each having a random, but unknown heading. A fraction of nodes are landmarks, meaning that they have self positioning capability by an external method such as GPS. Gaussian noise is added to each AOA estimation to simulate measurement errors. Gaussian distribution has the property that 95% of the samples lie within 1.96 standard deviations from the mean. What this means for angle measurements is that if the standard deviation of the noise is for example $\frac{\pi}{8}$, then 95% of the measurements will be in the interval $(-\frac{\pi}{4}; \frac{\pi}{4})$ of the true bearing, thus giving a total spread of $\frac{\pi}{2}$ for bearing measurements. Performance will be evaluated based on the accuracy of positioning for non-landmark nodes, accuracy of heading, and percentage of the regular nodes which succeed the solving for a position (coverage). All the results presented in this section are averaged from 100 runs with different randomly distributed landmark configurations over the same network. Due to the fact that the proposed algorithms provide different tradeoffs, in order to produce comparable coverage we ran *DV-Bearing* with a TTL of 5 and *DV-Radial* with a TTL of 4. In both cases the angle threshold was 0.35 ($\simeq 20°$). All performance graphs indicate the standard deviation in selected points.

[1] *isotropic* = having the same physical properties in all directions (connectivity, density, node degree, landmark distribution)

Positioning error (figure 11) is represented relative to the maximum communication range of a node. An error of 1.0 means that the position resulted from the positioning algorithm is one (maximum sized) radio hop away from its true position. For *DV-Bearing*, this position is obtained from the bearings to landmarks, applying the triangulation method mentioned in section II-B. For *DV-Radial,* position is obtained from the radials by solving a linear system. On the horizontal axis of the graphs the standard deviation of the measurement noise is varied from 0 to $\frac{\pi}{4}$, and the several curves on each graph correspond to different landmark ratios. A larger number of landmarks improves both accuracy and precision, by solving a larger system for each positioning problem. For reasonable errors *DV-Radial* provides better positioning accuracy, and exhibits less dependence on the percentage of landmarks.

Bearing error (figure 12) is the average error of the

Fig. 13. Heading error

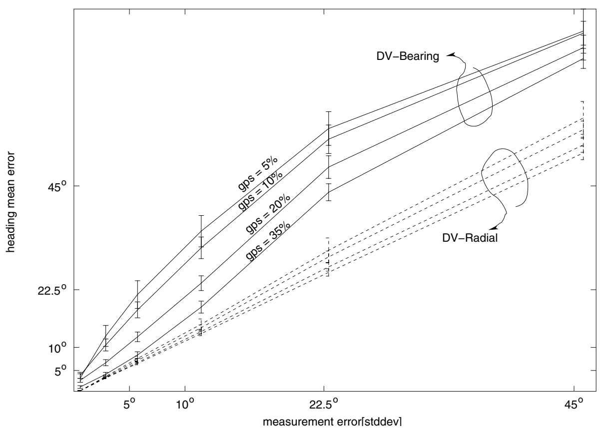

Fig. 14. Coverage

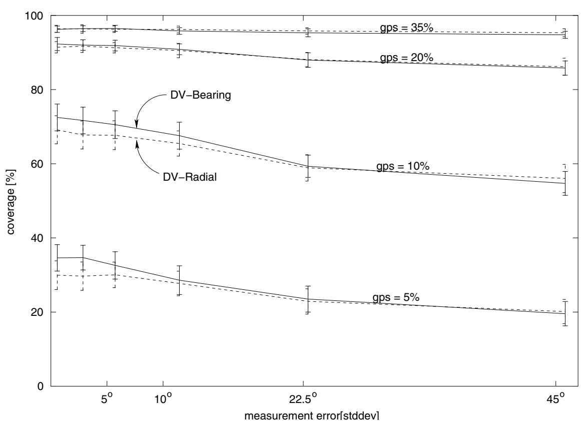

bearing to landmarks obtained by regular nodes after the orientation propagation phase stops. This is a primary measure of how the forwarding method compounds and propagates error. Because each landmark is treated independently, bearing errors are not affected by the number of landmarks available in the network. As expected, *DV-Radial* exhibits lower error, mainly because of the extra value that is forwarded.

Heading is the angle between nodes axis and the north, as would be given by a compass. Heading error is therefore the error in the absolute orientation averaged over all nodes. In our simulation, it is obtained by each node after estimating a position. Heading error (figure 13) is about double the bearing error, which is consistent with the results presented in [11].

Coverage (figure 14) represents the percentage of non landmark nodes which are able to resolve for a position. The reasons for which a node doesn't get a position are: fewer than three landmarks accumulated (due to propagation errors), collinear or co-circular landmarks, or numerical instability in the system solving. We aimed for similar coverage for the two algorithms in order to compare the other performance metrics. Even if positioning is theoretically possible with two landmarks for *DV-Radial* and with three landmarks for *DV-Bearing*, in practice, due to angle errors compounding, a much higher number of landmarks might be needed.

The main observations to draw from simulations are the following: accuracy can be traded off for coverage by tuning the TTL and the threshold value. The TTL tradeoff is also between energy and coverage, as its reduction would lead to less energy spending but also to less coverage. Positions obtained are usable for applications such as geodesic routing, as it is showed in [5], with errors of similar scale. Bearing errors follow closely the measurement noise, but they can be further decreased using more sophisticated correction methods.

In order to evaluate the accuracy of positions and orientations for a realistic application, we devised a simple example in which a mobile traverses a fixed network and is sensed by nodes within a certain distance. Nodes are initially deployed randomly, with a fraction of them (20%) having the self positioning capability. After running the APS algorithm to infer their position and orientation, the nodes sensing the mobile report their position and the direction in which the mobile was observed. At a central location, reports from various nodes are aggregated to produce an estimate position of the mobile. Since both positions and directions reported by nodes are based on APS produced positions/orientations, and therefore affected by errors, and because there may be more than two reporting nodes, the estimate position of the mobile is obtained by solving an over-determined linear system, in order to minimize the square error. In figure 15, the original trajectory is shown with a dashed line, and the restored one with a solid line. Standard deviations are indicated for each sample point. While more complicated data fusion/prediction techniques (such as Kalman filters) may be used here to improve the estimated trajectory, the purpose of the example is merely to quantify the APS produced error in the position and orientation of the nodes, with no additional processing. The network used (fig 10) was an isotropic topology with 100 nodes, mean degree 8.18, 20 nodes of which have self positioning capability. The measurement error considered was white gaussian noise with a standard deviation of 0.08 radians, which is about double the error of 5^o achieved by the AOA nodes realized by the Cricket compass project[6]. The algorithm used to infer position and orientation is *DV-Bearing*, which trades off some accuracy in order to work with less signaling and fewer capabilities (no compasses anywhere in the network). We assumed that the sensing distance is equal to communication radius, so that for each point we get about 6 or 7 readings. The sensing angle error is assumed to be 0, so that all the errors in the restored trajectory quantify the errors in our positioning algorithm

Fig. 15. tracking example isotropic topology

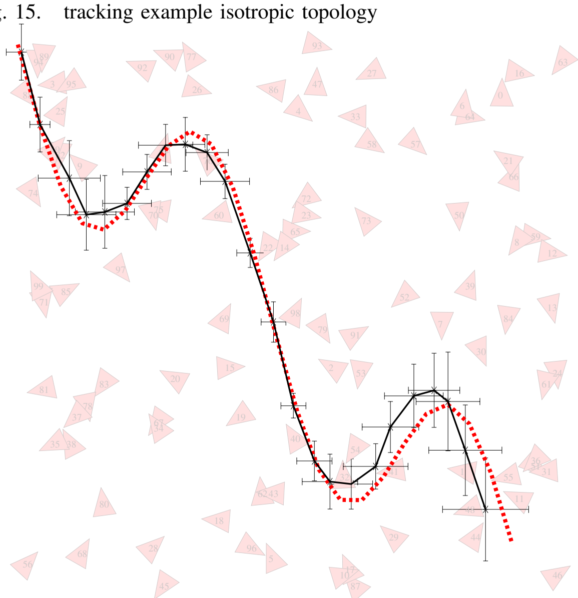

(*DV-Bearing*). It is interesting to note that estimations in the middle of the network are much more accurate that the ones at the edge (and this was verified with various other trajectories). The main cause for this is that an observation at the edge is obtained from angles which are clustered together in a small zone of the trigonometric circle - for example, a corner estimation would have all the angles in one single quadrant. In fact this is true about positions obtained by both algorithms. This would suggest that this class of algorithms (positioning, orientation, tracking) would run better when the border of the network is reduced in size, or is directly supported by preferential landmark placement.

VI. NODE MOBILITY

Our current simulation of APS only considers static topologies. While highly mobile topologies, usually associated with ad hoc networks, would require a great deal of communication to maintain up to date location, we envision ad hoc topologies that do not change often, such as sensor networks, and indoor or outdoor temporary infrastructures, like disposable networks. APS aims to keep a low signaling complexity in the event network topology changes slightly. When a node moves, it will be able to get distance vector updates from its new neighbors and triangulate to get its new position, therefore communication remains localized to nodes that are actually mobile. Not even moving landmarks would cause a communication surge in our approach because the only things that identify a landmark are its coordinates. In fact, a moving landmark would provide *more information* to the positioning algorithm, as the new position of the landmark acts as a new landmark for both mobile and fixed nodes. With a mobile "landmark", we can envision a case when a single, fly-over GPS enabled node is in fact enough to initialize an entire static network. Subsequent mobility of the network is supported as long as a sufficient fraction of nodes remains fixed at any one time to serve updates for the mobile nodes. While APS would perform well for limited mobility, it is very likely that its DV nature would incur high signalling costs in highly mobile scenarios. Drawing from experience in ad hoc routing, we may infer that an on-demand positioning scheme would be more appropriate for these cases.

An avenue that is explored extensively in mobile robotics research, involves usage of accelerometers and gyroscopes. Situations may arise when either a node doesn't have enough neighbors to get sufficient orientation readings, or the node wishes to stay in an inactive state for security or power conservation reasons. In this cases dead reckoning could be used to infer an estimate of current position based on the last triangulated position. This capability is given by accelerometers, which can provide relative positioning after a double integration of acceleration readings. Heading can be inferred in a similar manner when gyroscopes is available.

VII. FUTURE WORK AND CONCLUSIONS

Besides the extensions to mobility, already mentioned, future development of the project will be in the direction of improving the positioning quality by using error estimation and multimodal sensing.

An error estimation method proposed in [14] involves transmitting of the error estimation together with DV data. A node performing the orientation estimation described in section III-A, would also compute the estimated error of the newly computed orientation and forward it along. By increasing the signalling overhead, the final triangulation method has the possibility of using weights for each landmark. In case of range based APS, this provided a considerable reduction in positioning error. We are still investigating the adaptation of this error method for angle propagation.

Multimodal sensing can enhance the performance of positioning algorithms. AOA and ranging, possibly enhanced with compasses and accelerometers, have the possibility to provide better positioning than any of them taken separately. Both AOA and ranging are or can be currently achieved using common hardware - time difference of arrival (TDoA), based on ultrasound transmitters/receivers. Not requiring additional hardware makes multimodal based sensing a viable approach for positioning, which we plan to explore in the future.

To conclude, the method we proposed infers position and orientation in an ad hoc network where nodes can measure angle of arrival (AOA) from communication with their immediate neighbors. The assumption is that all nodes have AOA capability and only a fraction have self positioning capability. Two algorithms were proposed *DV-Bearing* and *DV-Radial,* each providing different signaling-accuracy-coverage-capabilities tradeoffs. The advantages of the method are that it provides absolute coordinates and absolute orientation, that it works well for disconnected networks, and doesn't require any additional infrastructure. What makes the algorithm scalable to very large networks is that the communication protocol is localized. Simulations showed that resulted positions have an accuracy comparable to the radio range between nodes, and resulted orientations are usable for navigational or tracking purposes.

REFERENCES

[1] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister, "System architecture directions for networked sensors," in *ASPLOS-IX*, Cambridge, MA, November 2000.

[2] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *ACM MOBICOM*, August 2000, boston, MA.

[3] A. Savvides, C.-C. Han, and M. Srivastava, "Dynamic fine-grained localization in ad-hoc networks of sensors," in *ACM MOBICOM*, 2001, rome, Italy.

[4] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *INFOCOM*. IEEE, March 2000, tel Aviv, Israel.

[5] D. Niculescu and B. Nath, "Ad hoc positioning system (APS)," in *GLOBECOM*. IEEE, November 2001, san Antonio.

[6] N. Priyantha, A. Miu, H. Balakrishnan, and S. Teller, "The cricket compass for context-aware mobile applications," in *6th ACM MOBICOM*, July 2001, rome, Italy.

[7] G. Finn, "Routing and addressing problems in large metropolitan-scale internetworks," University of Southern California, Tech. Rep. ISI Research Report ISI/RR-87-180, March 1987.

[8] J. C. Navas and T. Imielinski, "Geographic addressing and routing," in *MobiCom'97*, September 26-30 1997, budapest, Hungary.

[9] Y.-B. Ko and N. H. Vaidya, "Location-aided routing (lar) in mobile ad hoc networks," in *MobiCom'98*, October 1998.

[10] N. Bulusu, J. Heidemann, and D. Estrin, "GPS-less low cost outdoor localization for very small devices," in *IEEE Personal Communications Magazine*, ser. Special Issue on Smart Spaces and Environments, October 2000.

[11] M. Betke and L. Gurvitis, "Mobile robot localization using landmarks," in *IEEE International Conference on Robotics and Automation*, vol. 2, May 1994, pp. 135–142.

[12] B. Parkinson and J. Spilker, *Global Positioning System: Theory and Application*. American Institute of Astronautics and Aeronautics, 1996.

[13] B. Krishnamachari, S. Wicker, and R. Bejar, "Phase transition phenomena in wireless ad-hoc networks," in *GLOBECOM*, November 2001, san Antonio, TX.

[14] D. Niculescu and B. Nath, "DV based positioning in ad hoc networks," *Telecommunication Systems, Baltzer*, vol. 1, 2003, to appear.