

Localized Positioning in Ad Hoc Networks

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

Abstract—Position centric approaches, such as Cartesian routing, geographic routing, and the recently proposed trajectory based forwarding (TBF), address scalability issues in large ad hoc networks by using Euclidean space as a complementary name space. These approaches require that nodes know their position in a common coordinate system. While a GPS receiver in each node would be ideal, in many cases an approximation algorithm is necessary for networks with only a few GPS enabled nodes. These algorithms however require collaboration of large portions of the network, thus imposing an overhead for nodes which do not need positioning, or are mobile. We propose Local Positioning System (LPS), a method that makes use of local node capabilities – angle of arrival, range estimations, compasses and accelerometers, in order to internally position only the groups of nodes involved in particular conversations. Localized positioning enables position centric uses, like discovery, flooding and routing in networks where global positioning is not available.

Index Terms—trajectory based forwarding, positioning, local positioning system

I. INTRODUCTION

The main features of new ad hoc networks include large number of unattended nodes with varying capabilities, 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, easy to deploy, and most importantly easy to maintain. 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. In all these networks, both basic network operations (routing, forwarding), and higher level functions (multicast, resource discovery) impose a tradeoff between communication overhead and infrastructure support. For example, a network may have high powered basestations to hold large routing tables, or it has to use flooding to discover routes on demand; it may have a GPS [1] receiver in each node, or it has to spend some energy running a positioning algorithm [2]. Another tradeoff is encountered in route management - either proactively maintain routes to all possible destinations, or reactively discover them when needed. Both approaches prove better than the other one under different mobility and communication conditions.

The scalability problem of very large ad hoc networks of small nodes, such as sensor networks, is that traditional routing

schemes are either table based, or source based, both of which incur high costs on the network. A table based routing scheme can require a high amount of memory at each node, depending on the number of possible destinations. It can require a large amount of signaling to maintain the table in case of mobility. A source based, on demand, scheme shifts the scalability problem to the discovery process, based on flooding, and to the packet overhead, sensitive to path length. A class of routing schemes that addresses these problems is using the Euclidean space as a complementary namespace for routing [3], [4], [5]. These are position centric approaches, because the position of the destination functions as a delivery address, thus eliminating both the need of routing tables at forwarding nodes, and the need for complete path indication in each packet.

Position centric approaches normally assume that positions of nodes are known in the network, as node positions are used for both naming and forwarding decisions. In this paper we argue that Cartesian routing, TBF, and the applications they enable would work even in a network in which positions are not globally known, provided that some localized sensing of neighboring nodes is available. The central idea of localized positioning is that only nodes maintaining active communications need to be positioned with respect to each other. Positioning can be relative, that is, each group of communicating nodes maintains a different coordinate system for the purpose of that communication. A group usually means two end points and the intermediate nodes in the case of routing, but may also comprise of nodes involved in TBF based discovery, or flooding. A point to multi-point communication group for example, can use a coordinate system belonging to the source, provided that the intermediate nodes and the receivers are all registered to that coordinate system.

The necessary node functionality to support localized positioning is also required by most positioning algorithms, and includes the ability to measure ranges to neighbors, angle of arrival (AoA), or orientation of the node with respect to the North. Ranging is usually provided by estimating the distance to a neighbor by measuring the strength of the radio signal from that neighbor, or by TDoA (Time Difference of Arrival), a method that sometimes employs radio and ultrasound signals. Angle of arrival (AoA) is a method to estimate the direction from which a neighbor is sending data. It can be implemented either as an antenna array, or as a combination of radio and ultrasound receivers. Digital compasses may be used in combination with AoA measurement to provide more information about a node's orientation. All these methods have been previously used to provide or extend positioning

capabilities in ad hoc networks [6], [7], [2], [8], [9], [10].

If positions are not available to all nodes in a consistent fashion, due to lack or obstruction of positioning infrastructure, such as GPS satellites, the only option left is to run a distributed positioning algorithm. These algorithms however, come with their own disadvantages - they may either require separate infrastructure, like in the case of Cricket [7], requiring ultrasound beacons, or may require a preprocessing phase in which large portions of the network collaborates, thus imposing a high communication cost on all nodes, every time the network changes topology, like in the case of APS [2].

What we propose, instead of positioning all the nodes in the network, which may be either globally inconsistent, due to mobility, or expensive, due to high signaling, is to locally position only the nodes involved in point to point communication, in a coordinate system relative to the originator of the packet. This enables some applications of TBF – flooding, discovery, or source-sink communication in networks in which global positioning is not available, or necessary.

The main sections of the paper are organized as follows: the next section reviews related work, section III is a short review of trajectory based forwarding, section IV presents Local Positioning System (LPS), our method to position only the nodes involved in communication (forwarding the trajectory). Section V discusses simulation results, and we summarize with some concluding remarks in section VI.

II. RELATED WORK

Cartesian routing [3] assumes that when all nodes know their positions, and the position of the destination is known, intermediate nodes can forward using greedy decisions, for example choosing the node that provides most progress towards destination. Geocasting [4] sends messages to all nodes contained in a designated polygon, but was designed for IP networks in which nodes are GPS enabled. A recently proposed scheme, trajectory based forwarding [11], generalizes the Cartesian scheme by allowing packets to travel along arbitrary trajectories. Global, precise locations like those available from GPS [1] are highly desirable for all these methods. When GPS cannot provide full coverage, extension or replacement systems have been designed. RADAR [6] is a positioning system that maps many signal strength readings at different locations of a building and stores them in a signal strength readings database. A location query provides the readings to its visible beacons, thus choosing a most likely location from the database. The system has good performance, but has the disadvantage of requiring the beacon infrastructure and manual mapping of the building. Cricket [7], [9] is an MIT project employing radio and ultrasound beacons in order to implement ranging and orientation capabilities for roaming mobiles. APS [2], [10] and AhLOS [8] are positioning algorithms which do not need in advance setting of infrastructure, but require a small percentage of the nodes in the ad hoc network to have self positioning capability. The schemes extend the positioning capability to the rest of the nodes by the means of network wide collaboration algorithms. Self Positioning Algorithm [12] proposes relative positioning, in which the coordinate system

is established by a reference group of nodes. Their approach inspired our present work on LPS, by suggesting that adjacent nodes can establish local coordinate systems and register them.

In order to use locations as a naming space, a locations service is necessary to translate node addresses in coordinates. GLS (Geographic Location Service) [13] implements a naming service that allows node centric applications to run on top of Cartesian, geographic or trajectory based routing. A source can find the coordinates of the destination node from the location service and then use geographic or cartesian routing to route to that destination. While the location service is not only useful, but necessary in a global positioning network, it is still a research issue what similar mechanism would be necessary in a network in which communication is based on relative coordinate systems.

III. TBF REVIEW

Trajectory Based Forwarding (TBF) [11], [5], is a paradigm which aims to provide solutions that require neither infrastructure support, nor communication overhead for route maintenance. The idea is to embed the trajectory in each packet and let the intermediate nodes take the forwarding decisions. TBF is a generalization of source based routing [14] and Cartesian forwarding [3]. Like in source based routing, the path is indicated by the source, but without actually specifying all the intermediate nodes. Like in Cartesian forwarding, decisions taken at each node are greedy, but are not based on distance to destination - the measure is the distance to the desired trajectory.

TBF combines the best of the two methods: packets follow a trajectory established at the source, but each forwarding node takes a greedy decision to infer the next hop based on local position information, while the overhead of representing the trajectory does not depend on path length. In a network where node positions are known, the packet may be forwarded to the neighbor that is geographically closest to the desired trajectory indicated by the source node. If the destination node is known, the trajectory followed by the packet might be a line, and the method reduces to cartesian forwarding. In the general case, however, one can envision a larger array of applications (fig 1), including flooding (a), disjoint (b) and braided (c) multi-paths, discovery (d), geocasting (e) and multicasting (f).


IV. LOCAL POSITIONING SYSTEM (LPS)

Localized positioning can be used in any node centric network, when a communicating group agrees on a common coordinate system. We describe LPS in the context of trajectory based forwarding because TBF naturally creates communication groups of nodes which are touched by the trajectory. The trajectory may be a curve describing a path, or a tree, describing a distribution path.

We extend an idea proposed in [12], and develop a method for nodes to use some capabilities (ranging, AoA, compasses) to establish local coordinate systems in which all immediate neighbors are placed. It is then possible to register all these coordinate systems with the coordinate system of the source of the packet. Local Positioning System (LPS) is a method

Fig. 1

TRAJECTORY BASED FLOODING, ROUTING AND DISCOVERY



to achieve positioning **only for the nodes along the trajectory**, with minor increase in communication cost, as if all node positions were known. Instead, each node touched by trajectory spends some computation to position itself in the coordinate system of the source of the packet. The tradeoff achieved by the method is to spend computation in order to save communication. The computation spent for registration is more economical energy-wise compared to the communication spent by a network wide positioning algorithm (for current hardware, such as the Berkeley mote [15]).

A. Node capabilities

The network model considered is a large scale, dense, ad hoc set of 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 and symmetrical, these networks are modeled as fixed radius random graphs, or unit graphs.

We present several methods that would enable trajectory based forwarding, based on relative positioning that use ranging and angle of arrival (AoA) measurements.

Range based estimation of distance between two nodes has been previously used for positioning purposes [2], [8], [6], even with the high measurement error involved. In figure 2, node A would have estimations of distances to its neighbors AB and AC , but could also find BC after communicating with B or C . In most implementations, ranging is achieved either by using an estimate from the strength of the signal (unreliable), or using time difference of arrival (TDOA).

In AoA approach, each node in the 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. After the deployment, the axis of the node has an arbitrary, unknown

Fig. 2

NODES CAPABILITIES: MEASURING ANGLES, RANGES, ORIENTATION





Fig. 3

LOCAL COORDINATE SYSTEMS



heading, represented in figure 2 by a thick black arrow. When interacting with two neighbors, 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 \hat{ac} and \hat{ab} , and has the possibility of inferring one angle of the triangle, $\widehat{CAB} = \hat{ac} - \hat{ab}$. For consistency all angles are assumed to be measured in trigonometric direction. 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 has been developed at MIT by the Cricket Compass project [9]. Its principle of operation is based on both time difference of arrival (TDOA) and phase difference of arrival. Other node capabilities that might be available in small form factors include accelerometers and compasses. The accelerometer’s main use is to indicate pose, while compass indicates the absolute orientation of each node.

B. Local coordinate systems

Denote with \mathcal{C}_i the coordinate system of node i . The source node of a trajectory has the coordinate system \mathcal{C}_0 . In figure 3, if node A is able to measure distances to all its neighbors, via a ranging method, it can compute the sides and the angles for all triangles created with all pairs of neighbors which are also neighbors of each other. This would enable A to place itself in $0, 0$ of its local coordinate system \mathcal{C}_a , and all its immediate neighbors at positions that satisfy all the range requirements known to A , ranges indicated by continuous lines in the figure. To establish its coordinate system, A randomly chooses E as an indicator for its x^a axis and F as an indicator for its y^a axis. All the other neighbors are then placed in this coordinate system such that all the range readings are respected. Most nodes can be unambiguously positioned if they are neighbors with two previously positioned nodes. Node G for example is positioned immediately after E and F are positioned by choosing between two possibilities. It can be placed where it is on the figure, or on the position of G_1 . But if G were G_1 , it would have been close enough to E to have a direct link, so this possibility is eliminated. In the same fashion, all nodes around $A(0, 0)$ are positioned. Node B sets up a similar local coordinate system \mathcal{C}_b , by initiating its axes using neighbors C and D . Since a condition of global connectivity in a random ad hoc network is for the average degree to be at least 6 [16], in higher density networks most nodes will succeed in placing all their neighbors in their local coordinate systems.

In the example in figure 3, a translation and a rotation are enough to overlap the local system of B , \mathcal{C}_b over the local system of A , \mathcal{C}_a . In the case of D , however, the first two neighbors randomly chosen as axes indicators produce a coordinate system that cannot be overlapped over \mathcal{C}_a using only translation and rotation. In this case, due to D 's original, localized and independent choice of axes, a mirroring transformation will also be necessary. The main cause for these mirroring transformations is the fact that ranging information does not provide a sense of direction (at least in the first phase, before the local coordinate systems are set).

A similar procedure to establish a local coordinate system is possible when using angle of arrival measurements. In the same figure 3, node A knows all angles to neighbors reported against its main axis (see section IV-A), and therefore has already an established set of axes. In triangle AFE for example, A knows all angles but no sides. Assuming that the range $AE = 1$, it can then find the sizes AF and FE which can then be propagated to all other triangles in order to get coordinates for all points. This coordinate system, \mathcal{C}_a can be registered with \mathcal{C}_b using a scale transformation (in addition to translation and rotation), because B will probably choose a different edge as a unit for its local coordinate system. If a compass is available in each node, all the reference systems have parallel axes, and the rotation is not necessary to achieve coordinate system alignment. A mirroring transformation will occur when using AoA only in the case in which a node is deployed upside down.

It is possible to use both AoA and ranging in creating local coordinate systems, possibly enhanced with local compasses.

TABLE I
HARDWARE CAPABILITIES AND TRANSFORMATIONS

Capability	Transformations
Range	T, R, M
AoA	T, R, S, (M)
AoA+Compass	T, S, (M)
AoA+Range	T, R, (M)
AoA+Range+Compass	T, (M)

Table I indicates all the possible combinations of node capabilities, and the transformations involved in the alignment process (T=translation, R=rotation, S=scaling, M=mirroring). When mirroring is indicated in parenthesis, it can only happen as a result of a node being deployed upside down, not from the randomness in starting the local coordinate system. When only ranging is used, mirroring is possible regardless of the pose of the node, depending on the nodes chosen as indicators for local axes. In all the other cases, since AoA is assumed to report angles in the same (trigonometric) direction for all nodes, mirroring between two local coordinate system appears only when one node is flipped, situation which can be robustly detected by a digital accelerometer. The general transformation matrix to translate a point from \mathcal{C}_b to \mathcal{C}_a is:

$$M_{b,a} = \begin{bmatrix} sr_1 & sr_2 & t_x \\ sr_3 & sr_4 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

where $\{r_*\}$ is the rotation transformation and possibly mirroring, s the scaling factor, $\{t_*\}$ the translation. $\{r_*\}$ is an orthonormal matrix with the following properties:

$$\begin{cases} |r_1| = |r_4| \\ |r_2| = |r_3| \\ \det\{r_*\} = -1 \text{ for mirroring, } 1 \text{ otherwise} \end{cases}$$

If a transformation is not present, its parameters will be set to neutral values to preserve the properties of M . For example, when ranging is available and no scaling is involved, $s = 1$, when compasses are available and there is no rotation, $\{r_*\} = I_2$ (the identity matrix).

If for example, $C_a = [x_a^C \ y_a^C \ 1]^T$ designates the position of C in \mathcal{C}_a , and C_b its position in \mathcal{C}_b , then

$$C_a = M_{b,a} C_b \quad (1)$$

C. Registration

The registration between two local coordinate systems is the process that computes a transformation matrix which will overlap any point from one coordinate system to the other. The input to this process are points for which the coordinates are known in both coordinate systems with some accuracy. If perfect ranging were used in creating the local coordinate systems, the registration would produce a rigid transformation.

If two nodes, such as A and B in figure 3 wish to agree their coordinate systems, they must implement the registration procedure. They reach this agreement by using points that are common in the two coordinate systems, in our example

A, B, C and D . In practice, since ranging does not provide perfect distances and AoA perfect angles, the two coordinate systems will not overlap in all points if we restrict M to the type of transformation described above. The error in registration is defined as the sum of squared distances between corresponding points after registration. If B has the two sets of coordinates of nodes A, B, C and D in both coordinate systems, it computes the transformation matrix $M_{b,a}$ that minimizes the error:

$$\sum \|X_a - M_{b,a}X_b\|^2, X = A, B, C, D \quad (2)$$

Fortunately, there is a closed form solution to this optimization, described in [17], that has a running time linear in the number of common points. The solution for the rotation R involves computing of eigenvectors and eigenvalues for a 2×2 matrix (for a 2D problem). The scaling transformation s is obtained as the square root of the ratio of the variances of the two sets around their centroids. The translation t is the difference of centroid of set a and scaled and rotated centroid of set b .

D. Forwarding using LPS

The aim for LPS is to make forwarding along the trajectory similar to the procedure followed in a network where node positions are available. The key idea is that the only nodes that are positioned are the ones involved in forwarding along the trajectory. Positioning is done in a hop by hop fashion, in the coordinate system chosen by the initiating node - the source of the packet. The forwarding procedure works with a node selecting the next hop based on the proximity to the desired trajectory, or any of the other possible policies. In figure 3, the ideal trajectory is shown as a thick dashed arrow. Assume that A is already registered with \mathcal{C}_0 , that is, it has computed the matrix $M_{a,0}$. If the next node to be selected along the trajectory is B , it will receive from node A the position of common nodes in \mathcal{C}_0 , $\{A, B, C, D\}_0$, so that node B , can register its own coordinate system to \mathcal{C}_0 by solving the optimization problem (2). Once $M_{b,0}$ is obtained at B , all the neighbors of B are evaluated in \mathcal{C}_0 , by using transformation (1). Node B is then able to select one of its own neighbors that is closer to the trajectory, in order to continue the process.

What is in fact achieved by LPS is the registration of all coordinate systems of visited nodes to the coordinate system of the initiating node, which achieves positioning of all these nodes in the coordinate system of the source. This positioning system has a number of advantages:

- it is localized to the nodes actually involved in communication. Unlike a network wide positioning algorithm, such as [2], [8], which involves collaboration and coordination of a large number of nodes, LPS involves only the nodes “touched” by the desired trajectory;
- the size of packet only depends on the number of common neighbors two consecutive nodes have, and this is upper bounded by the maximum degree of the nodes. In the forwarding packet, the equation of the trajectory is also included, but this is a fixed size, so the packet does not increase with the length of the trajectory;

- can make use of any of the two previously used localization capabilities - AoA (possibly enhanced with local compasses and accelerometers), or range estimation;
- can be used for flooding, discovery, and position centric routing between static end-points, because all may rely on relative positioning.

E. Reducing the overhead for static networks

Registration at system i produces a matrix

$$M_{i,0} = \text{register}(\mathcal{C}_i, \mathcal{C}_0)$$

which can translate local coordinates $(x, y)_i$ in \mathcal{C}_0 coordinates:

$$(x, y)_0 = M_{i,0}(x, y)_i$$

The amount of packet overhead to support LPS comprises:

- 1) $(X(t), Y(t))_0$ - encoded either using universal tables to reduce overhead, or reverse Polish notation, to allow maximum flexibility. These functions are meaningful in \mathcal{C}_0 only;
- 2) t_{crt} - the point of the curve closest to the current node;
- 3) coordinates of common nodes - $(A, B, C, D)_0$ in coordinate system \mathcal{C}_0 ;

If TBF is used with a global coordinate system, registration is not required, and only (1) and (2) are necessary. If Cartesian routing is used, the line between the current node and the destination describes the trajectory, therefore only the (3) is an overhead (unless when global coordinates are used, in which case there is no packet overhead).

The amount of CPU overhead for forwarding at node i :

- 1) Register \mathcal{C}_i to \mathcal{C}_0 :

$$M_{i,0} = \text{register}(\{A, B, C, D\}_i, \{A, B, C, D\}_0)$$

- 2) Translate all neighbors of node i in \mathcal{C}_0 :

$$H_0 = M_{i,0}H_i, \forall H \text{ neighbor of } i$$

- 3) Evaluate $X(t), Y(t)$ in the interval $[t_{crt}, t_{crt} + \Delta]$, to find the next hop, in \mathcal{C}_0 ;

If TBF is used with a global coordinate system, the only overhead item is (3).

By examining the two categories of overhead - computation and communication, we can see that they are dominated by registration related items: (3) for packet overhead, (1) and (2) for CPU overhead. Registration is what makes all the nodes operate in a common coordinate system, the core of LPS. The only problem is that it claims a per packet cost in terms of both CPU and packet overhead, a cost caused by the absence of a global coordinate system.

In static networks, this problem can be addressed by caching. The procedure of successively aligning coordinate systems as the trajectory goes is important for discovery purposes, or for the cases when topology changes. When the network is static, per packet registration of coordinate systems for each packet may be avoided by running a preprocessing procedure in which all nodes only register once with each neighbor, and cache the transformation matrix for that neighbor for all subsequent communication, thus reducing both CPU

and packet overhead. Instead of registering with the origin of communication, which is only known at forwarding time, register with the neighbor:

$$M_{i,0} = M_{i,i-1}M_{i-1,0}$$

$$M_{i,0} = \prod_{j=i}^1 M_{j,j-1}$$

and use the fact that the registration matrix $M_{i,0}$ is in fact the product of all the registration matrices from the source of the packet to the current node. This means that in fact node i only needs to register with the neighbor labeled $i-1$ ($M_{i,i-1}$), in order to handle all communication forwarded from $i-1$, which is independent of packet's coordinate system. What is specific to a particular source with coordinate system \mathcal{C}_0 at node i is $M_{i-1,0}$, which can be forwarded from the node $i-1$, together with the trajectory. The matrix is usually small in size, depending on the hardware used, see table I at page 4. Translation requires two numbers t_x and t_y , rotation is an angle, and the scale is also a number. These values replace the coordinates of common nodes (item (3) of packet overhead), whose number grows with the node degree, and achieve the registration step (item (1) of CPU overhead) with a simple matrix multiplication. Therefore the per packet cost imposed by LPS to support TBF in a static network without global positioning is reduced to four numbers as packet overhead. The per packet computation overhead imposed by LPS is one matrix multiplication per neighbor. The preprocessing cost involves establishing local coordinate systems, registering individually with each neighbor, and storing the transformation matrices for subsequent communication.


V. SIMULATION

We simulated an isotropic¹ map (average degree=9.5), with 200 nodes each having a random, but unknown heading. In order to evaluate LPS for the purposes of forwarding, routing, and discovery we tested various linear trajectories, and used four metrics to compare the performance of forwarding in the presence of ranging errors and angular errors. For each of 30 pairs of source-destination nodes situated near the sides of the map (North-South and East-West), we compared the trajectory obtained when ranging or angle measurement is imperfect with the trajectory obtained when these measurement are perfect. In the latter case all nodes register to their globally coherent positions. Therefore, the “perfect trajectory” is the one obtained with perfect locations, which is still different from the ideal trajectory, as described by the parametric curves. Note that globally coherent positions all differ from true GPS positions by that same transformation.


- **Deviation** of a trajectory is computed as the mean distance between a point on the obtained trajectory and the closest point on the perfect trajectory. The value of this deviation is then normalized to the maximum hop size to be expressed as a hop count.

¹isotropic = having the same physical properties in all directions (connectivity, density, node degree)


Fig. 4
RANGE, RANGE + AOA, RANGE + AOA + COMPASS




(a) Deviation from 'perfect' trajectory



(b) Path length - relative to 'perfect' trajectory



(c) Minimum distance to destination



(d) Probability of reaching the destination


- **Path length** is the number of hops achieved by a trajectory, relative to the perfect trajectory. The forwarding process may be stopped when the trajectory exits the map, or when packets cannot be forwarded anymore due to local aberrations in registration, obstacles, faulty or inexistent local coordinate systems.
- **Minimum distance to destination** shows how close a trajectory gets to the desired destination.
- **Probability of reaching destination** expresses the chances of the trajectory passing within one hop of the desired target. In this set of experiments, when the trajectory reaches the destination, or an immediate neighbor of it, the forwarding is stopped.

For the angle of arrival, we model normally distributed errors whose standard deviation is shown on the horizontal axis. Corresponding to an error of $\pm 5^\circ$, achieved by the Cricket compass project corresponds a standard deviation of 2.5° on our graphs. For ranging measurement, we assume a linear error model where the actual measurement is within a fixed fraction of the true range (0-100%).


Simulations are divided in two groups, based on the node capabilities used by LPS. The first group (figure 4) compares ranging (labeled as 'SS') with ranging and AoA, possibly enhanced with a compass. The AoA deviation is fixed in these experiments to 0.15 (about 8.59°). The experiments reveal that ranging is less robust in forwarding the coordinate system, being greatly enhanced by the presence of angular measurements, even with high errors. The reason for this difference stems from the way local coordinate systems are built. Using ranging with errors, the coordinate systems are prone to false mirroring, large aberrations produced by working with small or obtuse angles, bias depending on the nodes initializing the coordinate system. We experimented with various thresholds to limit the shape of triangles resulted from ranging - eliminating very obtuse or very acute angles from being used in inferring local coordinates. Such thresholds provide a tradeoff between the amount of possible misleading information that is accepted in the registration and the number of nodes which successfully get a local coordinate system. For this set of experiments, we used an angular threshold of 0.6 ($\sim 36^\circ$), thus eliminating from the inference the triangles with smaller angles, at the price of dropping around 3% of the nodes which weren't able to establish local coordinate systems. When AoA is available, the local coordinate system building relies on the inherent order in which angles are reported, completely eliminating false mirroring, and reducing range aberrations. Range based measurement shows a maximum in the deviation (figure 4a), explained by the sharp decrease in the path length (figure 4b) following the maximum, which produces shorter paths with a relatively good start.

The second group of experiments (figure 5) studies the behavior of angular measurements, possibly enhanced with a compass. The absence of ranging in these experiments brings the need for a scaling transformation in the registration process. On the horizontal axis of all these graphs, the standard deviation of the angular measurement error is indicated. The deviation when using simple AoA (figure 5a) shows a decreasing trend that is caused by the paths getting shorter with


Fig. 5
AOA, AOA + COMPASS




(a) Deviation from 'perfect' trajectory



(b) Path length - relative to perfect trajectory

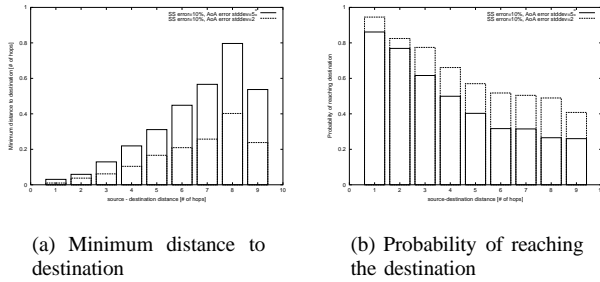


(c) Minimum distance to destination



(d) Probability of reaching the destination

Fig. 6
ERRORS COMPOUND WITH DISTANCE



increasing errors, but still having a good start, which yields a low deviation. The same type of behavior was encountered when a single type of sensing was employed (ranging), and is a symptom indicating short paths. For the region where the deviation is high, path length can be even longer than the perfect path, because, once it misses the destination, it can wander for 1-2 more hops before it reaches the edge of the map. Both groups of experiments show that a compass in each node brings a substantial improvement in all metrics. One of the main causes for this improvement is that rotation is eliminated from the registration process ($\{r_*\} = I_2$).

These two sets of simulations showed that multi-modal sensing (AoA + ranging) performs much better than its components taken separately, under the same error conditions. A compass attached to each node brings significant improvement in all the metrics considered.

In order to evaluate the amount of degradation that occurs in the iterative registration process, we conducted a set of experiments in which source and destination nodes are chosen randomly on the map, with uniform probabilities for X and Y coordinates. In figures below, minimum distance and probability of reaching the destination are plotted as a function of the initial distance between source and destination, normalized by the maximum hop size. LPS employed ranging with an error of 10% and AoA measurements with deviations of 2° and 5° . As expected (figure 6), the error is compounding with distance yielding decreasing probabilities of reaching the destination and increasing distances to destination. An interesting side effect is that for far apart source - destination pairs, which will usually be placed on the opposite sides of the map, the minimum distance between the trajectory and destination decreases, as the trajectory may shortly wander left or right, searching for a possible continuation of the map, and thus reducing the distance to the destination.

The main conclusions revealed by the simulations are that, although error compounds with distance, it can be countered efficiently by using additional hardware available in small form factors. An accelerometer present in each node could detect node flipping, eliminating false mirroring, while a digital compass could eliminate rotations from registrations process, when angular measurements are used. With AoA error range in the range $2^\circ - 5^\circ$, which has already been realized by other projects [9], LPS maintains a low deviation, making

it usable for discovery purposes, and reaches the destination with very high probability, making it usable for routing.

VI. CONCLUSIONS

We presented Local Positioning System (LPS), a method to achieve localized positioning in ad hoc networks where certain capabilities (AoA, ranging, compasses, accelerometers) are present. LPS requires only the collaboration of nodes involved in a communication group, be it a trajectory, or a distribution tree. LPS does not incur the communication overhead of a network wide positioning algorithm by registering only the coordinate systems that are in actual communication.

When the network is static, per packet registration of coordinate systems for each packet may be avoided by running a preprocessing procedure in which all nodes only register once with each neighbor, and cache the transformation matrix for that neighbor for all subsequent communication, thus reducing both CPU and packet overhead.

Simulations showed that trajectories positioned by LPS have a low deviation, making them usable for discovery purposes, and reach the destination with very high probability, making them usable for routing.

REFERENCES

- [1] B. Parkinson and J. Spilker, *Global Positioning System: Theory and Application*. American Institute of Astronautics and Aeronautics, 1996.
- [2] D. Niculescu and B. Nath, "Ad hoc positioning system (APS)," in *GLOBECOM*, San Antonio, November 2001.
- [3] 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.
- [4] J. C. Navas and T. Imielinski, "Geographic addressing and routing," in *ACM MOBICOM*, Budapest, Hungary, September 26-30 1997.
- [5] D. Niculescu and B. Nath, "Trajectory based forwarding and its applications," Rutgers University, Tech. Rep. DCS-TR-488, May 2002, DATAMAN Lab, Department of Computer Science.
- [6] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *INFOCOM*, Tel Aviv, Israel, March 2000.
- [7] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *ACM MOBICOM*, Boston, MA, August 2000.
- [8] A. Savvides, C.-C. Han, and M. Srivastava, "Dynamic fine-grained localization in ad-hoc networks of sensors," in *ACM MOBICOM*, 2001.
- [9] N. Priyantha, A. Miu, H. Balakrishnan, and S. Teller, "The cricket compass for context-aware mobile applications," in *ACM MOBICOM*, Rome, Italy, July 2001.
- [10] D. Niculescu and B. Nath, "Ad hoc positioning system (APS) using AoA," in *INFOCOM*, San Francisco, CA, April 2003.
- [11] B. Nath and D. Niculescu, "Routing on a curve," in *HOTNETS I*, Princeton, NJ, October 28-29 2002.
- [12] S. Capkun, M. Hamdi, and J.-P. Hubaux, "GPS-free positioning in mobile ad-hoc networks," *Cluster Computing*, vol. 5, no. 2, April 2002.
- [13] J. L. et al, "A scalable location service for geographic ad hoc routing," in *ACM MOBICOM*, August 2000.
- [14] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," *Mobile Computing*, vol. 353, 1996.
- [15] 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.
- [16] B. Krishnamachari, S. Wicker, and R. Bejar, "Phase transition phenomena in wireless ad-hoc networks," in *GLOBECOM*, San Antonio, TX, November 2001.
- [17] B. Horn, H. M. Hilden, and S. Negahdaripour, "Closed form solution of absolute orientation using orthonormal matrices," in *Journal of the Optical Society of America*, July 1998, vol. 5, no. 7, pp. 1127-1135.