

A Robustness Analysis of Multi-hop Ranging-based Localization Approximations

Kamin Whitehouse
Computer Science
UC Berkeley

kamin@cs.berkeley.edu

David Culler
Computer Science
UC Berkeley

culler@cs.berkeley.edu

ABSTRACT

In this study, we implement six ranging-based localization algorithms from the literature and evaluate them in simulations that employ real-world ultrasound ranging data. We find that small variations in the ranging model can lead to large variations in localization error. We analyze each algorithm to identify how implicit assumptions may be violated by empirical ranging data and why this changes the behavior of the algorithm.

Categories and Subject Descriptors

I.6.5 Modeling Methodologies, C.3 Embedded Systems

General Terms

Algorithms, Performance

Keywords

Localization, Noise Models, Ranging, Sensor Networks, Simulation, Statistical Emulation

1. INTRODUCTION

Ranging-based sensor localization is the problem of identifying the locations of sensor nodes given estimates of the distances between them, known as *range* estimates. There are currently a large number of ranging-based localization algorithms in the literature, each of which makes a different geometric approximation: some algorithms assume that a network can be decomposed and localized as several sub-networks; other algorithms assume that range estimates can be added together to create longer range estimates; other algorithms assume that multi-dimensional coordinates can be projected onto a two-dimensional space. Each of these approximations greatly simplifies the sensor localization problem and has been validated in simulation with the *Noisy Disk* ranging model. However very few of them have been tested or demonstrated with real-world ranging techniques

such as ultrasonic time of flight or radio signal strength. This raises the following important question:

Question *Do approximation algorithms make assumptions that may be violated by real-world ranging techniques?*

In this study, we implement six ranging-based localization algorithms from the literature and analyze each of them for implicit assumptions about the ranging model. Our scientific approach is to first evaluate each algorithm with the Noisy Disk ranging model. Then, we add a new ranging irregularity X to our model (eg. non-Gaussian noise), and evaluate the algorithm again. A comparison of the two results verifies one of two possible hypotheses:

H_0 *Irregularity X does not violate an assumption of the algorithm, and localization error in both trials will be the same.*

H_1 *Irregularity X does violate an assumption of the algorithm, and localization error in both trials will be different.*

We repeat this procedure for several different ranging irregularities. Once we have established *which* assumptions of an approximation algorithm are violated by empirical ranging data, we perform a deeper analysis and discussion into *how* and *why* these violations affect each of the six approximation algorithms.

Several of these algorithms have been designed to be robust to ranging *noise*, ie. they are designed such that small variations in range estimates do not cause large variations in localization error. In contrast, this study evaluates an algorithm's robustness to the ranging *model*; it tests whether small variations in the model cause large variations in localization error. In many ways, robustness to the ranging model is more important than robustness to ranging noise; it has been well established that real-world ranging characteristics are unlike those of the Noisy Disk model [14] and, furthermore, every deployment will have distinct, idiosyncratic, and anomalous ranging features due to different hardware, environmental sensitivities, and signal analysis techniques. A localization algorithm that directly relies on artifacts of a specific parameterization or theoretical model is therefore unlikely to be useful in practice. The results of this study argue that robustness to the ranging model and to ranging irregularities should be a standard metric of evaluation for localization algorithms in the future. To our knowledge, this is the first study to evaluate most of these algorithms with a ranging model other than the Noisy Disk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN'06, April 19–21, 2006, Nashville, Tennessee, USA.
Copyright 2006 ACM 1-59593-334-4/06/0004 ...\$5.00.

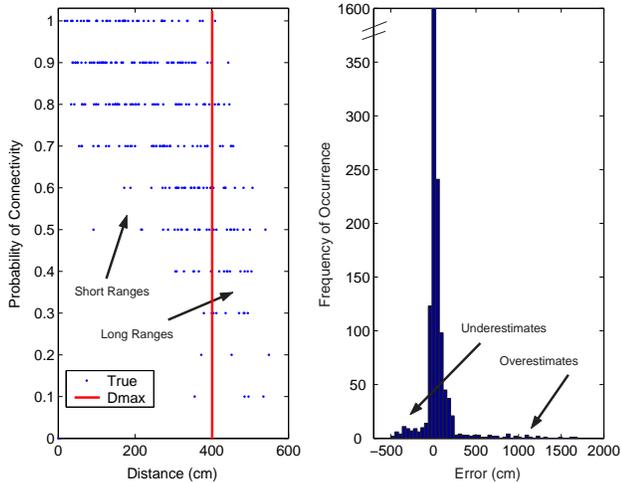


Figure 1: The (a) connectivity and (b) noise distribution of empirical ultrasound ranging data exhibit all four types of ranging irregularities.

2. EXPERIMENTAL SETUP

For the purposes of this study, we define sensor field localization in terms of n nodes in a two-dimensional plane, the first m of which are *anchor nodes* which have known locations in a global coordinate system. We focus on ranging-based localization, in which each node i can obtain a distance estimate to its neighbor j that is some function of the true distance between them.

$$\hat{d}_{ij} = f(d_{ij}) \quad (1)$$

Many pairs of nodes i and j , especially those with large distances between them, will not obtain a valid distance estimate, which we call a *range failure* and denote as $\hat{d}_{ij} = \emptyset$. Thus, the nodes and the distance estimates between them form a graph $\mathcal{G} = (V, E)$, where $|V| = n$ and $e_{ij} \in E = \hat{d}_{ij}$. Ranging-based sensor field localization is the problem of deriving the positions of the $n - m$ unlocalized nodes from the graph \mathcal{G} .

2.1 Ranging Models & Irregularities

A ranging model is the function $f(d_{ij})$ from Equation 1. The *Noisy Disk* is the idealized model of ranging in which each node obtains a range estimate with Gaussian noise σ to all neighbors within a maximum range d_{max} . The Noisy Disk is the *de facto* standard model for evaluating or comparing most localization algorithms in the literature. Range estimates can be generated according to this model in simulation through the following function:

$$\hat{d}_{ij} = \begin{cases} \mathcal{N}(d_{ij}, \sigma) & d_{ij} \leq d_{max} \\ \emptyset & d_{ij} > d_{max} \end{cases} \quad (2)$$

In a previous study, we used the Noisy Disk to predict the localization error of real deployments using both ultrasound and radio signal strength ranging. Once we actually performed the deployments, however, we found that the empirical localization error was dramatically different from that predicted [14]. This motivated us to identify four different types of empirical ranging irregularities:

Extreme overestimates: an excess of range estimates that are longer than the true distance by more than two standard deviations

Extreme underestimates: an excess of range estimates that are shorter than the true distance by more than two standard deviations

Long-range proficiency: the existence of range estimates between nodes farther than nominal range d_{max}

Short-range deficiency: the existence of range failures between nodes closer than nominal range d_{max}

Each of these irregularities can be seen in our empirical ultrasound data shown in Figure 1. Figure 1.b shows a histogram of ranging error between several hundred pairs of nodes at different distances. Each pair uses the filtered value from 10 independent ranging trials. The extreme underestimates occur when the ultrasound transducer detects an ultrasound signal before it actually arrives, perhaps due to environmental noise or calibration problems. The extreme overestimates occur when the transducer does not detect the signal until long after it actually arrives, perhaps due to multi-path effects. Figure 1.a shows, for these same pairs, the probability of successfully obtaining a range estimate on each of 10 trials at different distances. The short-range deficiency is evident in that many pairs that are closer than d_{max} do not in fact obtain a ranging estimate with some probability. This may be caused by irregular environmental attenuation or because our noise filter has removed highly variable readings. A long-range proficiency is also evident in that pairs that are farther than d_{max} do obtain a range estimate with some non-zero probability, perhaps caused by an amplifying environmental pathway.

In this study, we evaluate each algorithm with five different ranging models, each of which incorporates one ranging irregularity more than the previous one:

Model 1 (Noisy Disk): No irregularities

Model 2: Model 1 + Extreme overestimates

Model 3: Model 2 + Extreme underestimates

Model 4: Model 3 + Long-range proficiency

Model 5: Model 3 + Short-range deficiency

We use the empirical ultrasound ranging data from Figure 1 to generate these ranging irregularities in simulation with a technique that we call *Statistical Emulation*, which uses statistical sampling techniques to employ empirical data directly in simulation. We previously validated this technique by comparing it to real localization deployments and showing that it accurately predicts localization error [14]. Complete details about the Statistical Emulation algorithm are available in that study.

In all of our models, the Noisy Disk parameters σ and d_{max} are derived directly from the ultrasound data in order to make the three models as comparable as possible. We fit a mixture of Gaussians to the ultrasound data and choose the parameter σ with the highest likelihood. This value is more representative of the nominal data than fitting a single Gaussian because it is not affected by noise irregularities. The maximum range d_{max} is set to a value

| Single-hop | Multi-hop | | |
|--|--|---|----------|
| | Centralized | Distributed | |
| <ul style="list-style-type: none"> • RADAR • Active Bats • Cricket • MoteTrack | <ul style="list-style-type: none"> • Max Likelihood • Convex Opt | <ul style="list-style-type: none"> • DV-distance • Bounding Box | Absolute |
| | <ul style="list-style-type: none"> • MDS-Map | <ul style="list-style-type: none"> • GPS-free • Robust Quads • MDS-Map(P) | |

Figure 2: This taxonomy shows the six algorithms evaluated in this paper in bold.

that would achieve the same average degree as the empirical data. We call this value the *effective range* of the empirical data, which can be calculated as $r_{eff} = \frac{\sqrt{P}}{\Pi}$ where

$$P = \int_{r=0}^{r=d_{max}} \Pi * r^2 * p(r) \quad (3)$$

and $p(r)$ is the probability of successfully obtaining a range estimate at distance r . Nodes using the Unit Disk model of connectivity with parameter $d_{max} = r_{eff}$ should have the same number of neighbors as nodes using the empirical ultrasound data.

2.2 Algorithms and Topology

The six algorithms we evaluate are shown in Figure 2.2 in terms of a traditional taxonomy of localization algorithms, and are each described in more detail when analyzed in later sections of the paper. Several of these algorithms use a form of *refinement* [7, 9, 5, 10] in which location estimates are improved through either coordinate or gradient ascent. We do not include refinement algorithms in our analysis because they can be applied to the results of any localization algorithm and do not affect their approximations or assumptions.

We evaluate all six algorithms on a simulated 18m x 18m square topology with 100 nodes and four anchors on each side. Nodes are placed in a grid with Gaussian noise added to the grid positions to avoid exhibiting artifacts of the network partitions which are likely in a completely random topology or of the strict rigidity of a grid, neither of which are representative of the canonical deployment. Regardless of the ranging model used, all networks have an average degree of 9, which has been shown to be a transition point between high and low density networks [4]. The only exception to this rule are trials with Model 4, for which it was impossible to hold both d_{max} and average network degree constant. These trials have higher network degree of 12 because the model incorporates long-range proficiency *and* all nodes closer than d_{max} are connected.

This experiment only measures a single point along the dimensions of network size, anchor density, topology, etc. The chosen topology is sufficiently representative of the canonical sensor field deployment, however, because smaller topologies

| Rank | Ideal | Empirical |
|------|------------------------------|-----------------------------|
| 1 | MDS-Map DV-Distance | MDS-Map |
| 2 | MDS-Map(P) | DV-Distance Bounding Box |
| 3 | Robust Quads Bounding Box | MDS-Map(P) |
| 4 | GPS-Free | Robust Quads GPS-Free |

Figure 3: Algorithm rankings differ when using the idealized ranging model and empirical ranging data.

exhibit predominantly edge effects (nearly all nodes are near an edge) and larger deployments can be subdivided into a network of this type by placing anchor nodes appropriately throughout the network. The purpose of this study is not to explore the effects of network size and anchor density, which has been done in other studies [4, 6, 8], but to explore the effects of ranging irregularities.

3. EXPERIMENTAL RESULTS

The median errors and the median percentage of nodes localized in 30 trials of each algorithm with all five ranging models are shown in Figure 4. We used a one-tailed t -Test with $\alpha = 0.05$ to formally compare adjacent models, and the statistically significant changes are indicated in the figure with *s. These changes allow us to accept one of the two hypotheses from Section 1 for each algorithm. For example, the Bounding Box and DV-Distance algorithms are both significantly affected by Models 3 and 4, but not Models 1 and 5. Therefore, we conclude that they are sensitive (H_1) to extreme underestimates and long-range proficiency and not sensitive (H_0) to extreme overestimates and short-range deficiency. We analyze the results of each algorithm in more detail in the following six sections.

The most surprising result of this study is the extremely high sensitivity of all algorithms to small changes in the ranging model. Simulation with a theoretical model is never an exact replica of reality and is never expected to produce exactly the same algorithmic response as empirical noise. However, we do typically expect simulation to be:

1. **indicative:** results should be within a constant *fudge factor* of empirical results
2. **decisive:** an algorithm that performs better in simulation should also perform better in reality

The results of our experiment do not meet either of these expectations. Localization error increases by less than 70% for some algorithms and over 800% for others. We used a one-tailed t -Test with $\alpha = 0.05$ to derive a statistically significant ranking of all algorithms, and the resultant orderings are very different when using Models 1 and 5, as shown in Table 3. This indicates that simulation using the Noisy Disk ranging model is neither indicative nor decisive, meaning that it has almost no value when trying to design, tune, and deploy a localization algorithm in the real world. This motivates our more detailed analysis and discussion in the next six sections about *why* each algorithm is so sensitive to ranging irregularities.

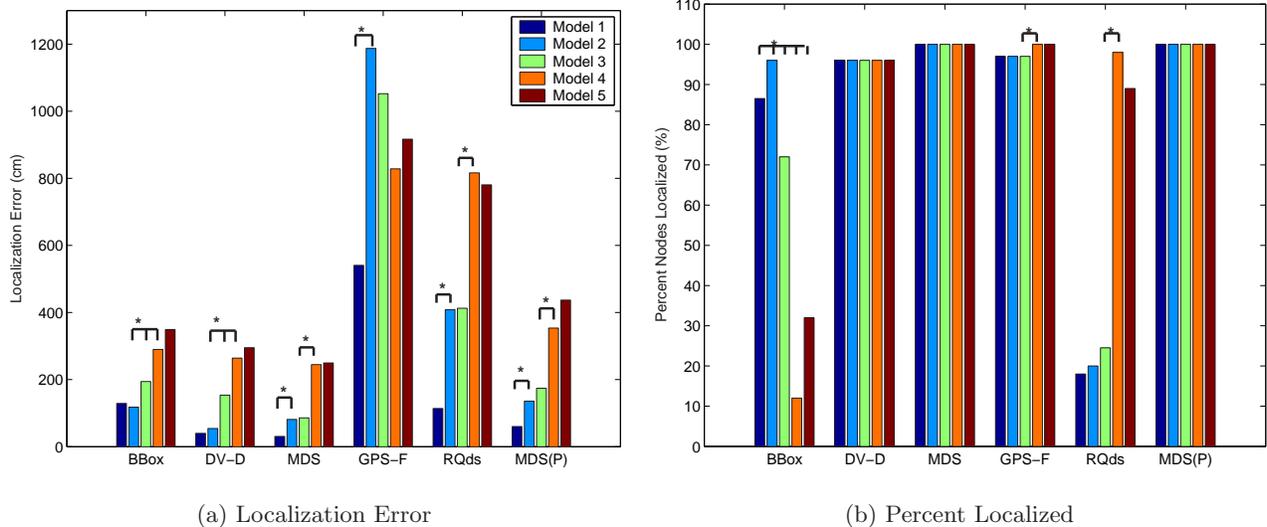


Figure 4: Each algorithm along the x -axis is subject to five ranging models. (a) shows the median localization error and (b) shows the median percentage of nodes localized over 30 trials. Statistically significant changes are indicated with *'s.

4. BOUNDING BOX

The Bounding Box algorithm [9, 12] uses the length of edge $e_{ij} \in E$ to constrain the unknown coordinates of node i in terms of the known coordinates of anchor node j . These constraints take the following form:

$$\begin{aligned} x_j - e_{ij} < x_i < x_j + e_{ij} \\ y_j - e_{ij} < y_i < y_j + e_{ij} \end{aligned}$$

These are very loose constraints which only require that a node be *within* a certain distance from an anchor node, not *at* a certain distance. Furthermore, the constraints are placed on the x and y coordinates independently, so the union of constraints from all anchor nodes defines a box, not a circle, as depicted in the left side of Figure 5.

Bounding Box makes two approximations: 1) The location of each node is approximated to be the center of the box defined by the union of all constraints 2) all ranging failures $e_{ij} = \emptyset$ can be approximated by shortest path distances sp_{ij} through the graph \mathcal{G} . In other words, if i does not have a range estimate directly to j but both have range estimates to a third node k , the edge distance may be estimated to be $e_{ij} = sp_{ij} = \hat{d}_{ik} + \hat{d}_{kj}$. The shortest path distances can easily be calculated in the network using a distance vector algorithm similar to those used in routing.

As shown in Figure 4(a), localization error for the Bounding Box algorithm significantly increases when subject to extreme underestimates and long-range proficiency. Simultaneously, the percentage of nodes that the algorithm is able to localize drops significantly. However, the algorithm is largely impervious to extreme overestimates and short-range deficiency. We can explain this phenomenon through a deeper analysis of how the shortest path approximation is affected by noise and connectivity irregularities.

Intuitively, shortest path distances are always longer than the true distance because of their *zig-zag* nature, as illustrated in Figure 6(a). Shortest paths straighten out as

the network density increases and should asymptotically approach the true distance as density goes to infinity. However, any shortest path algorithm will preferentially use underestimated ranging estimates and avoids overestimated range estimates; given a choice between two similar paths, the algorithm will choose the one with more underestimates, as illustrated in Figure 6(b). Therefore, the positive and negative errors in range estimates do not necessarily cancel out; shortest path algorithms are highly sensitive to underestimates and are largely impervious to overestimates.

For these reasons, the extreme underestimates in our empirical ranging data can cause disproportionate errors in shortest path estimates. They effectively create shortcuts through the network and divert many shortest paths, causing widespread errors through a *wormhole* effect. This is illustrated in Figure 7(a), which shows that the introduction of ultrasound noise irregularities can cause 50-75% errors in shortest paths, even though the nominal empirical ranging error is only around 5-10%; the few range estimates that are extremely underestimated cause errors in a large number of shortest paths. The same graph shows that the impact of ranging noise increases with network density. This is because the distance vector algorithm has more reasonable alternatives to a shortest path in a very dense network than in a very sparse one and will therefore have more opportunities to use underestimates or avoid overestimates.

The effects of empirical noise are exacerbated by the effects of long-range proficiency. The shortest path algorithm prefers to use long ranges and largely ignores short ranges because long ranges tend to decrease the shortest path distance, as illustrated in Figure 6(c). Therefore, the fact that our empirical ultrasound data has more long ranges and less short ranges than the idealized model means that shortest path distances will be decreased further, even though network density remains the same. This effect is illustrated in Figure 7(a) where, for each density, the shortest path dis-

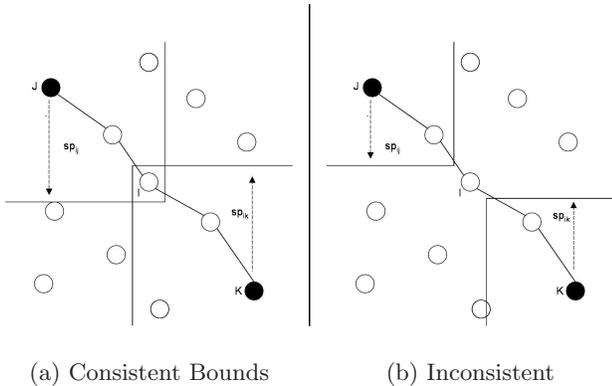


Figure 5: (a) Bounding Box constrains each node to be within its multi-hop distance to each anchor (b) sometimes resulting in inconsistent constraints.

tances are shorter when subject to both noise and connectivity irregularities than when subject to noise irregularities alone.

Thus, extreme underestimates and long-range proficiency combine to yield shortest path distances which are up to 50% shorter than the true distances. This is extremely detrimental to the Bounding Box algorithm when it results in *inconsistent* bounds, ie. the upper and lower bounds defined by two or more anchor nodes do not overlap. In this scenario, nodes cannot be localized, as seen in Figure 5(b). As noise and connectivity irregularities are introduced, the percentage of nodes localized by Bounding Box quickly drops to the extent that Bounding Box cannot localize most nodes when subject to Model 5, even at low densities. The shortest path errors at each density are strongly correlated with the percentage of nodes localized by bounding box, as shown by the two graphs in Figure 7.

5. DV-DISTANCE

Like the previous algorithm, DV-Distance [6] approximates the distance between a node i and an anchor node j to be the shortest path distance sp_{ij} . DV-Distance then uses this value to constrain the position of node i in terms of the position of each anchor node j with an equation of the following form:

$$0 = (x_i - x_j)^2 + (y_i - y_j)^2 - sp_{ij}^2$$

In contrast to Bounding Box, this strict equality relates both x and y coordinates in the same equation, forming a circular constraint. A system of at least three such equations can be linearized and solved with least squares for the coordinates x_i and y_i . In this way, DV-Distance uses the shortest path distances to reduce the multi-hop localization problem to single-hop localization and directly solve for node position.

Because DV-Distance makes the same shortest-path approximation as Bounding Box, it is susceptible to the same ranging irregularities: extreme underestimates and long-range proficiency. DV-distance solves directly for a point estimate of each node's position, however, and does not suffer from the problem of inconsistent bounds as Bounding Box does. Therefore, DV-Distance is always able to localize all non-anchor nodes in the network, as shown in Figure 4(b).

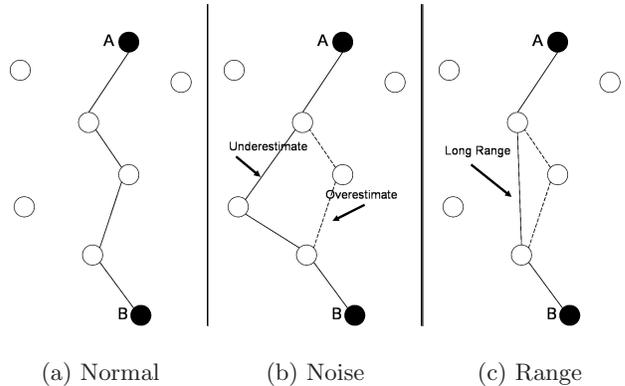


Figure 6: Theoretically, shortest paths zig-zag and should always be longer than the true distance. However, extreme underestimates and long-range proficiency combine to significantly shorten them.

To deal with noise in the shortest paths, DV-distance exploits the fact that each anchor node can calculate the ratio of the shortest path distance and the true distance to all other anchors. This ratio is transmit to the network as a *correction factor*, which can be used by neighboring nodes to adjust shortest path estimates before localization. The correction factor at anchor k for anchor j would be of the following form:

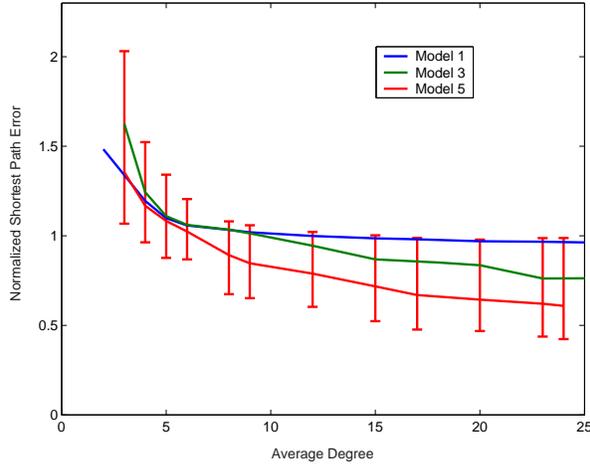
$$\frac{sp_{jk}}{\sqrt{(x_j - x_k)^2 + (y_j - y_k)^2}} \quad (4)$$

Anchor correction factors are intended to fix exactly the kind of systematic errors in shortest path distances that trouble Bounding Box. Indeed, Figure 5 shows that corrections factors cause the median shortest path distance error with Model 5 to be very similar to that achieved with Model 1. However, DV-Distance's localization error is affected by empirical ranging data almost as much as Bounding Box.

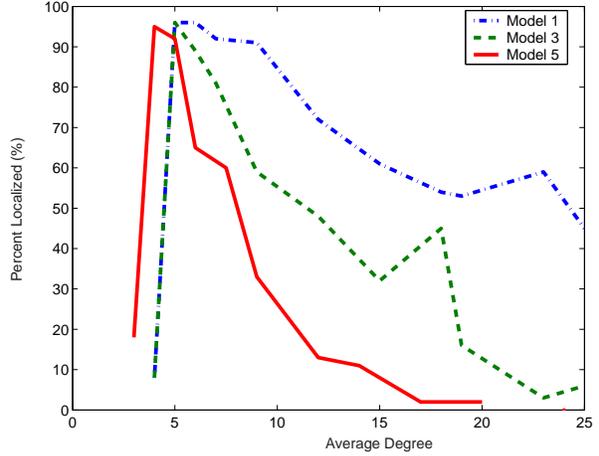
Correction factors are not very effective in the face of ranging irregularities because they do not remove the high *variance* in shortest path distances caused by Model 5. Figure 5 shows that corrected shortest path distances from Model 5 can be off by a factor of two both before and after correction factors. The reason for this high variance is that, as stated in Section 4, a small number of noise or connectivity irregularities cause errors in a large number of shortest path distances. However, they do not cause errors in all shortest paths, nor are all irregularities equally damaging. Therefore, some shortest path distances will be greatly affected by ranging irregularities while others will be unaffected. Anchor corrections apply to all shortest paths regardless, correcting any general bias in shortest path errors but not variance.

6. MDS-MAP

MDS is an analytical technique to derive n -dimensional positions of n objects given a complete similarity matrix D with the metric distances between them. MDS-map [11] is a ranging-based sensor localization algorithm that uses MDS by making two approximations: 1) all range failures $e_{ij} = \emptyset$ can be approximated by shortest path distances sp_{ij} to convert the incomplete matrix E into a complete similarity



(a) Shortest Path Errors



(b) Bounding Box: percentage localized

Figure 7: (a) shows that shortest paths get shorter as density increases and ranging irregularities are introduced. In (b), this leads to many nodes not being localized by the Bounding Box algorithm.

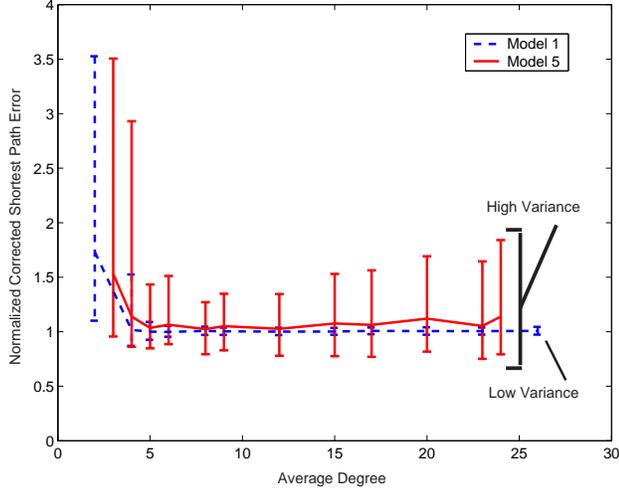


Figure 8: Correction factors reduce systematic bias, making the shortest path distances of Model 5 approximate those of Model 1. However, they do not remove the high variance that ranging irregularities cause in the shortest path algorithm.

matrix D 2) the locations of all nodes can be approximated by a 2-dimensional projection of the n -dimensional positions derived through MDS. This procedure requires the entire graph \mathcal{G} , so MDS-Map is a centralized algorithm.

Even though MDS-Map also uses the shortest path approximation, it is much more robust to underestimated ranges than Bounding Box and DV-Distance. One shortcoming of the previous two algorithms is that they only use edges e_{ij} between nodes and anchor nodes; DV-Distance uses one shortest path per anchor node, while Bounding Box only uses at most four shortest paths in total: those that define the highest and lowest bounds on its x and y coordinates. *MDS-Map* uses edges between all nodes simultaneously, dramatically increasing the number of constraints used to determine a node's location and reducing the ability of a few underestimates to have significant influence.

MDS-Map also shows a marked increase in sensitivity to extreme overestimates with respect to the Bounding Box and DV-distance, which were both impervious to them. This is because shortest paths of one or two hops are more likely to be affected by extreme overestimates than shortest paths of many hops, since the latter can more easily avoid the overestimate. Bounding Box and DV-Distance primarily use shortest paths to estimate long distances between nodes and anchor nodes. Because MDS-Map estimates the shortest path differences between *all* pairs of nodes, however, most of the pairs of nodes are relatively close together. While the first two algorithms therefore showed a general bias towards underestimated shortest paths, MDS-Map shows instead a much higher total variance, with many underestimated *and* many overestimated edge.

7. GPS-FREE

While the previous three algorithms are what we call *shortest path* approximations, the next three are what we call *patch and stitch* approximations. A patch and stitch algorithm forms a small patch or cluster around each node i con-

sisting of all neighbors $N_i \subset V$ where $j \in N_i \iff e_{ij} \neq \emptyset$. All nodes in each patch are localized relative to a local coordinate system and the overlap $N_{ij} = N_i \cap N_j$ between patches of neighboring nodes i and j can be used to derive a coordinate transform that will give the location of i relative to j . The relative locations of non-neighbor nodes i and j where $N_i \cap N_j = \emptyset$ can be calculated by cascading transforms from multiple overlapping patches between i and j . A global stitching order can be used to localize all nodes within the same coordinate system.

The patch localization algorithm used by GPS-free uses the center node i of the patch and two connected neighbor nodes $\{j, k \mid e_{ij}, e_{ik}, e_{jk} \in E\}$ to *bootstrap* a coordinate system by assigning the following coordinates:

$$\begin{aligned} (x_i, y_i) &= (0, 0) \\ (x_j, y_j) &= (\hat{d}_{ij}, 0) \\ (x_k, y_k) &= (\hat{d}_{ik} \cos(\gamma), \hat{d}_{ik} \sin(\gamma)) \end{aligned}$$

where

$$\gamma = \arccos\left(\frac{(\hat{d}_{ij})^2 + (\hat{d}_{ik})^2 - (\hat{d}_{jk})^2}{2(\hat{d}_{ij} * \hat{d}_{ik})}\right) \quad (5)$$

These three nodes become anchor nodes that define a local coordinate system. The other nodes in the patch are localized using *iterative multilateration* [9]: any node connected to at least three anchors is first localized and then used to localize other nodes. This process repeats until all possible nodes are localized. GPS-free chooses the two bootstrapped anchors j and k to maximize the total number of neighbors in the patch that can be localized through iterative localization. This criteria does not uniquely specify the pair, and in our implementation we randomly chose any pair that met this criteria.

Neither GPS-free nor Robust Quads, the next algorithm, specify a global stitching order so in our implementation we use the order specified by MDS-Map(P), the last algorithm. The set of stitched patches S is initialized to the largest patch $S = \operatorname{argmax}_i |N_i|$. The set of unstitched patches is set to be all other patches $\bar{S} = V - S$. At each step, the next patch to be stitched is determined to be the patch in \bar{S} that has the largest overlap with any patch in S

$$\operatorname{argmax}_j |N_i \cap N_j| \quad i \in S, j \in \bar{S} \quad (6)$$

The only significant change to the GPS-Free localization error is caused by extreme overestimates. However, this algorithm is actually sensitive to several ranging irregularities, but the effect is not statistically significant because the magnitude of the error has already approached that of random node placement. This extreme sensitivity to noise can be attributed to Iterative Multilateration, GPS-Free’s patch localization algorithm, which exhibits the same trend. Iterative Multilateration is highly sensitive to noise irregularities because each step in the process uses only a few range estimates, and subsequent steps build upon the results of earlier steps. While there are very few extremely underestimated or overestimated ranges in the data set, if one of them is used in an early stage of localization, the errors it causes will propagate to all other nodes in the cluster.

The more surprising fact about GPS-Free localization is that the percentage of localized nodes is actually *improved* when subjected to long-range proficiency and short-range

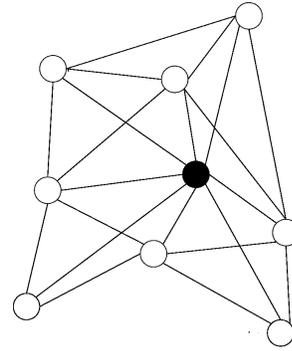


Figure 9: Using Robust Quads, the top six and the bottom three nodes in this cluster cannot be localized to a common coordinate system. At only 66% localized, this patch has a high probability of stopping the stitching process.

deficiency. These irregularities do not change the average degree of the network or the average number of nodes localized in each cluster. Instead, it makes the neighbors in each cluster slightly farther away from each other on average. This decreases the number of times that the intersection of two clusters N_{ij} is a co-linear set, causing a slight increase in the number of patches which can be stitched.

8. ROBUST QUADS

Robust Quads is a patch and stitch approximation that attempts to improve on GPS-Free by preventing some of the biggest errors in iterative multilateration. The robust quads algorithm defines a parameter of robustness θ which is usually set to 3σ , where σ is the nominal standard deviation of ranging noise. A triangle of neighbors i, j , and k is defined to be robust if $d \cos^2(\phi) > \theta$, where d is the shortest edge and ϕ is the smallest angle in Δijk . A clique of four nodes i, j, k, l is defined to be a *robust quad* if all triangles between these nodes form robust triangles. To localize a patch, this algorithm first finds all robust quads from the $\binom{|N_i|}{4}$ quadrilaterals in the patch. It forms an *overlap graph* $\mathcal{G}_o = (V_o, E_o)$ where each vertex is a robust quad $q \in V_o$ and vertices are connected iff the two quads have at least three nodes in common: $e_{qp} \in E_o \iff |q \cap p| = 3, q, p \in V_o$. Three nodes from one robust quad are used to bootstrap a coordinate system, as in GPS-free, and to localize the fourth node in that quad. Then, all other nodes in the patch are localized using iterative multilateration, with the order in which nodes are localized defined by a breadth-first search through the overlap graph. Because quads are fully connected, a quad that overlaps with a localized quad is guaranteed to have at most one unlocalized node that is connected to three localized nodes. Robust Quads does not specify how to choose the root of the breadth-first search. In our implementation, we choose any quad that contains the center node i of the patch. As with GPS-Free, we use the stitching order defined by MDS-Map(P).

Because Robust Quads tends to localize nodes using long ranges and avoids localizing nodes using short ranges, it is very sensitive to extreme overestimates and relatively robust to extreme underestimates. The dominant characteristic of Robust Quads, however, is that with Models 1-3 less than

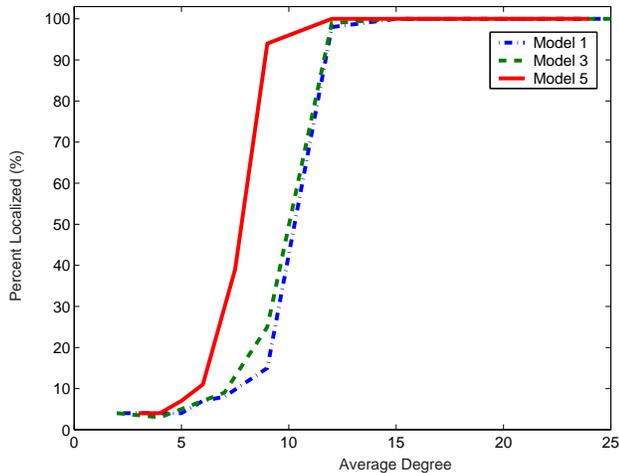


Figure 10: Robust Quads goes very quickly from localizing almost no nodes to localizing almost all nodes. Long-range proficiency triggers this phase transition at lower densities.

25% of the nodes can be localized to a global coordinate system in our topology. This is *not* true for the Robust Quads patch localization algorithm, which localizes on average 65% and 55%, respectively. Therefore, we can assume that the reason most nodes cannot be localized is because of a failure in the stitching algorithm.

Through inspection of several networks in which no more patches could be stitched together, we found that many of the patches on the fringe of the localized section were similar to the patch illustrated in Figure 9. This patch contains five robust quads, the top three of which overlap and the bottom two of which do not. Therefore, the top six nodes can be localized to a common coordinate system while the bottom three can not, i.e. 66% of the nodes could be localized. This causes a problem during the stitching process when only half of the cluster overlaps with a cluster that is localized to the global coordinate system; the cluster cannot propagate the coordinate system to the other side of the cluster, and stitching stops. In a network where most patches are localized to only 50-60%, the probability of each patch stopping the stitching process is high enough that only a small number of patches can be stitched.

Subjecting Robust Quads to long-range proficiency improves the percentage of nodes that can be localized to a global coordinate system. In the patch localization process, the percentage of nodes localized increases to about 80%. This is because a patch with more long ranges is more likely to have robust triangles. This is similar to the reason why GPS-Free localized more nodes when subjected to long-range proficiency. While a change from 50-60% to 80% is not very high, it is enough to trigger a *phase transition* in which the probability of each patch stopping the stitching process becomes low enough that most patches can be stitched. This phase transition is evident in Figure 8. With both Models 1 and 3, the transition occurs between average degrees of 10 and 12. When long-range proficiency is introduced, however, the transition occurs between degrees of 6 and 7.5. Long-range proficiency also causes the average localization error to significantly increase, approaching that of

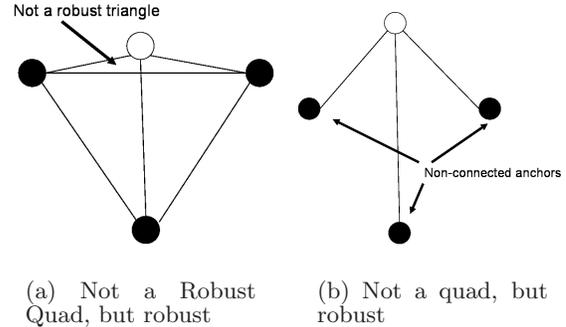


Figure 11: Robust Quads is overly-restrictive in the sense that it will not allow the three anchor nodes in either of these topologies to localize the fourth node, even though it could be performed robustly.

GPS-Free. This is presumably because the new nodes that are being localized are those with higher localization error.

It is worth mentioning that the patch in Figure 9 can be localized without any danger of *flip* or *discontinuous flex* ambiguities described in [5]. The patch is not completely localized, however, because the patch localization algorithm is too restrictive in that it does not allow multilateration even in cases where it can be performed robustly, as illustrated in Figure 11.

9. MDS-MAP(P)

MDS-Map(P) uses MDS-Map as a patch localization algorithm. The shortest paths between all nodes in patch i are calculated and combined with the range estimates to form a complete similarity matrix D_i , which is used to localize the nodes in the patch relative to each other. No anchor bootstrapping is required. The local coordinate systems are then stitched together to form a global coordinate system using the algorithm described earlier. The original MDS-Map(P) proposal suggests using patches of nodes that are more than one hop from the center of each patch. In our implementation, we use one-hop patches to make the algorithm comparable to the GPS-Free and Robust Quads algorithms.

The characteristics of MDS-Map(P) follow the same trends as those of MDS-Map: the algorithm is more sensitive to connectivity irregularities than to noise irregularities, although significantly affected by both. The main difference between the two algorithms is that MDS-Map(P) has consistently higher error, regardless of the ranging model used, because the stitching algorithm is amplifying the errors introduced during the patch localization process. By comparing the localization of MDS-Map and MDS-Map(P) on a network that required chains of up to 25 stitches, we can infer that the stitching process amplifies errors by a factor of two or less.

10. RELATED WORK

Two previous studies have also compared the error of localization algorithms using both idealized and empirical ranging. In one study, the authors hand calibrated their radios to resemble the idealized model as much as possible and still reported that, with three hop-count based al-

gorithms, empirical localization results were several times worse than predicted by simulation [13]. In our previous work, we identified how noise from both ultrasound and radio signal strength ranging affect one localization algorithm at multiple different densities [14]. To our knowledge, the current paper is the first to study the effects of empirical ranging characteristics on multiple ranging-based approximation algorithms, and to analyze how and why empirical ranging characteristics affect each approximation algorithm.

Previous work has also addressed concerns of empirical noise and connectivity irregularities in different ways. At least one study assumes a Gaussian distribution from non-Gaussian acoustic range errors [3] by use of the Central Limit Theorem, which states that the means of samples from a population will tend to be normally distributed around the population mean, regardless of the distribution that describes the sampled population. This strategy does not eliminate extreme under or overestimates, however, when a pair of nodes *consistently* gets highly erroneous readings due to multi-path, attenuators, or other factors that affected the empirical data set used in this study. Several studies in the networking literature have addressed non-disk like connectivity in radio communication and the effect of this on applications [2, 15, 1, 16]. To our knowledge, this paper is the first to study non-disk like connectivity in ranging, and to analyze how this violates tacit assumptions of existing ranging-based localization approximations.

11. CONCLUSIONS

In this paper, we find that small variations in ranging model can cause large variations in localization error for several algorithms. This motivates a deeper analysis, in which we identify several implicit assumptions and analyze why each algorithm is affected by different ranging irregularities.

The results of our study argue that robustness to the ranging model is an important metric of evaluation that should be used for localization algorithms in the future. Highly sensitive algorithms may perform well in simulation but will be unpredictable when deployed with real-world ranging techniques like ultrasound or radio signal strength or when deployed in a new environment with attenuators and multi-path effects, all of which produce irregular and idiosyncratic ranging characteristics.

The most important way to increase robustness is to evaluate algorithms using real-world ranging data through simulation techniques such as Statistical Emulation. A major cause of the lack of robustness we observed is that the Noisy Disk ranging model is used to design, evaluate, and compare almost all localization algorithms and is the basis for most mathematical proofs and analysis. This allows algorithms to rely on specific artifacts of this one parameterization. Furthermore, while our models may continually improve, parametric ranging models will never incorporate unexpected ranging irregularities for thorough robustness testing. Unfortunately, most techniques that do use real world noise, such as trace-based simulation, emulation, and pilot deployments, can only evaluate the algorithms in a single topology, limiting analysis in other ways. Statistical Emulation can both accurately capture properties of real world noise and evaluate algorithms in any topology. As we saw in this paper, it can also incorporate only certain

features of empirical data, providing the algorithm designer with a scientific approach to identifying algorithmic assumptions and sensitivities. The Statistical Emulation simulator we used, called Silhouette, and empirical data sets for both ultrasound and radio signal strength ranging are all publicly available.

12. REFERENCES

- [1] A. Cerpa, J. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical model of lossy links in wireless sensor networks. In *IPSN*, 2005.
- [2] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex Behavior at Scale: An Experimental Study of Low-Power Wireless Sensor Networks. UCLA Computer Science Technical Report UCLA/CSD-TR 02-0013.
- [3] Y. H. Hu and D. Li. Energy based collaborative source localization using acoustic micro-sensor array. In *IEEE Workshop on Multimedia Signal Processing*, 2002.
- [4] K. Langendoen and N. Raijers. Distributed localization in wireless sensor networks: a quantitative comparison. *Computer Networks*, 43(4):499–518, November 2003.
- [5] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. In *SenSys '04*, 2004.
- [6] D. Niculescu and B. Nath. Ad Hoc Positioning System (APS). In *IEEE GLOBECOM*, pages 2926–2931, 2001.
- [7] C. Savarese, J. M. Rabaey, and J. Beutel. Locationing in distributed ad-hoc wireless sensor networks. In *ICASSP 2001*, 2001.
- [8] A. Savvides, W. Garber, S. Adlakhia, R. Moses, and M. Srivastava. On the error characteristics of multihop node localization in ad-hoc sensor networks. In *IPSN*, 2003.
- [9] A. Savvides, H. Park, and M. B. Srivastava. The bits and flops of the n-hop multilateration primitive for node localization problems. In *WSNA*, 2002.
- [10] Y. Shang and W. Ruml. Improved MDS-based Localization. In *IEEE Infocom*, March 2004.
- [11] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz. Localization from mere connectivity. In *MobiHoc*, June 2003.
- [12] S. Simic. A distributed algorithm for localization in random wireless networks. unpublished manuscript, 2002.
- [13] R. Stoleru and J. A. Stankovic. Probability grid: A location estimation scheme for wireless sensor networks. In *SECON*, 2004.
- [14] K. Whitehouse, C. Karlof, A. Woo, F. Jiang, and D. Culler. The effects of ranging noise on multihop localization: an empirical study. In *IPSN*, 2005.
- [15] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *The First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.
- [16] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic. Impact of radio irregularity on wireless sensor networks. In *Mobisys*, 2004.