

THE EFFECTS OF RANGING NOISE ON MULTI-HOP LOCALIZATION: AN EMPIRICAL STUDY

Kamin Whitehouse, Chris Karlof, Alec Woo, Fred Jiang, and David Culler

{kamin,ckarlof,awoo,xjiang,culler}@cs.berkeley.edu
Computer Science Division
UC Berkeley

ABSTRACT

This paper presents an empirical study of how ranging affects multihop localization in wireless sensor networks. We compare real-world deployments with simulation and establish an objective metric with which to evaluate ranging models: if a model accurately predicts real-world localization performance, it accurately captures ranging characteristics. We evaluate a well-established ranging model called *Noisy Disk* and find that it yields optimistic predictions; real deployment error can be four to five times worse than simulation error. A systematic analysis reveals that *Noisy Disk* fails to capture important connectivity characteristics which dominate the effects of error on localization.

1. INTRODUCTION

Multihop localization in wireless sensor networks enables nodes to determine their locations without direct connectivity to nodes in known positions. Simulation is an important tool for evaluating multihop localization algorithms, but we have discovered that real-world performance is often much worse than predicted by simulation. This discrepancy is consistent with the anecdotal experience of many researchers in the area and the dearth of systematic comparisons in the literature. The observed *prediction gap* is presumably due to differences between theoretical noise models and empirical noise characteristics in the ranging measurements used by the localization algorithm. This paper presents an empirical evaluation of the *Noisy Disk* model, which is used almost universally to model ultrasound and radio signal strength, and identifies where and how it deviates from real-world characteristics. We use the prediction gap as a quantitative metric of evaluation: if a model accurately predicts real-world localization performance, it effectively captures empirical ranging characteristics.

We perform real-world localization deployments using both ultrasound and radio signal strength and show that the observed localization error is much worse than that predicted by the *Noisy Disk* model. We propose a new and more accurate method of simulation that uses statistical sampling techniques and empirical data in simulation. We systematically replace each component of the *Noisy Disk* model with increasingly accurate models to quantify each compo-

ment's contribution to the prediction gap. Our results indicate that both empirical noise and connectivity characteristics deviate from the *Noisy Disk* model, and we demonstrate these deviations have significant impact on multihop localization performance.

2. BACKGROUND

A basic building block of localization is *ranging*, the process of estimating the distance between a pair of nodes. Two common ranging technologies are radio signal strength (RSS) and ultrasonic time of flight (TOF), both of which introduce noise and uncertainty to localization. RSS techniques estimate the distance between two nodes by assuming a known rate of signal attenuation over distance and measuring the strength of the received RF signal. RSS is sensitive to channel noise, interference, attenuators and reflections, all of which have significant impact on signal amplitude. RSS also suffers from transmitter, receiver, and antenna variability. Ultrasonic TOF estimates distance by assuming a constant speed of sound and measuring the time it takes for an acoustic signal to travel between a pair of nodes. Because TOF relies on the speed of the signal instead of the magnitude, it is relatively robust to most sources of noise including attenuators and reflections; the line-of-sight signal should always arrive at the same time, though it may be stronger or weaker.

For theoretical analysis and simulation, ultrasound and radio signal strength are almost universally modeled with a *Noisy Disk* model, which has two components: noise and connectivity. The noise component indicates the distribution of the error between the measured distance and the actual distance (e.g., Gaussian, uniform). The connectivity component indicates the maximum distance d_{max} between two nodes at which a distance estimate can be obtained. For example, using Gaussian noise, the *Noisy Disk* defines the distance estimate \hat{d}_{ij} between nodes i and j in terms of the true distance d_{ij} as

$$\hat{d}_{ij} = \begin{cases} \mathcal{N}(d_{ij}, \sigma) & d_{ij} \leq d_{max} \\ \text{undefined} & \text{otherwise.} \end{cases} \quad (1)$$

The *Noisy Disk* model with no noise component (i.e., it only models the connectivity between nodes) is also known as

the Unit Disk model.

The Noisy Disk model is ubiquitous in localization research, but researchers generally acknowledge that noise is not perfectly Gaussian or uniform, and connectivity is not disk-like. Regardless, it is still a useful model of noisy ranging estimates. Theoretical analyses have successfully used the Noisy Disk model to mathematically derive the maximum likelihood solution to localization [1], lower bounds on localization error [2, 3], or specific properties about localization algorithms [4]. The Noisy Disk Model is more commonly used to evaluate and compare algorithms in simulation [5, 6, 7, 8, 9, 10, 11, 12]. Several projects collected empirical ultrasound data [13] or RSS data [14, 15] to derive realistic values for the parameters d_{max} and σ , which are then used in simulating the behavior of various localization algorithms. Other studies use these parameters for sensitivity analysis by, for example, measuring accuracy while varying d_{max} from 1.1 to 2.2 times the average node spacing and σ from 0 to 50% of d_{max} or similar values [13, 16, 17, 18]. Although the Noisy Disk model has been useful for evaluating and developing multihop localization algorithms, no study has verified that it accurately predicts the performance of real-world deployments.

There are two fundamentally different classes of localization algorithms: single hop and multihop. Single hop localization assumes all nodes in the network have direct ranging connectivity with a set of nodes in known positions, called *anchor nodes*. Several commercial and academic real-world systems using single hop localization have achieved accurate results [19].

The main drawback of single hop localization is the direct connectivity requirement between nodes and anchors. To remove this assumption, researchers have developed multihop localization algorithms. However, multihop localization introduces many new challenges. While single hop localization requires only local computation on each node, multihop localization requires long-distance information transfer and node collaboration. Furthermore, multi-hop localization requires evaluation at large scale, in contrast to single-hop localization for which the results of a single-cell deployment can be generalized to larger multi-cell deployments because each cell is roughly independent. Because of these challenges, multi-hop localization research is still mainly focused on theoretical analysis and simulation, with relatively few successful large scale deployments. In this paper, we focus exclusively on multi-hop localization because it relies much more heavily on high fidelity ranging models to understand error propagation and for use in theoretical analysis and simulation. A survey of multihop localization algorithms can be found here [18].

3. DEPLOYMENT SETUP

We performed several medium-scale localization deployments with our localization system and present three of them in this paper. The first is a 49 node network over a 13x13m asphalt area localized using ultrasound. The others are 25



Fig. 1. Ultrasound Deployment. 49 nodes deployed in a random grid pattern in a parking lot were localized using ultrasound. The ultrasonic transducer and reflective cone are visible above the node.

and 49 node networks over a 50x50m grassy area localized using RSS. These three deployments were chosen in part because they represent the canonical multihop deployments for which many localization algorithms have been designed and which most localization simulations try to emulate. They also provide particular insight into the nature of the Noisy Disk model, as we will see later. Here we present the ranging and localization systems we used for these deployments, which builds upon and improves some of the best hardware designs and algorithms from several other systems to create a unified system that is specially tailored to this localization problem.

Our study utilizes state of the art ranging capabilities for a large 2D field of sensor nodes. For our RSS deployments, we chose a low-power radio from several that have been characterized for use with RSS ranging. An early study showed the RFIDeas badge system to yield 5m range and errors near 2m at 2m range [20], and later studies, including our own, characterized low-power ASK radios such as the RFM DR3000 and the RFM TR1000 [20, 21, 22] to yield about 1.5m error at 3m distances and up to 6m error at 6m distances, even in near-ideal conditions. In our deployments, we use the newer Chipcon CC1000 FSK radio, which was shown in a recent single-hop localization study to provide RSS fidelity similar to that of more sophisticated 802.11 radios [23]. Our own characterizations show that in near-ideal conditions and with low transmission power, the radio has a standard deviation in RSS readings that translates to about 2m ranging error at the maximum range of about 20m, after calibration.

Our ultrasound hardware combines and improves ideas from several ultrasound implementations. Our ultrasonic transducer circuitry is derived from that of the Medusa node

[13], except we add a switchable circuit so that a single transducer could be used to both transmit and receive. Our nodes measure ultrasonic time of flight by transmitting the acoustic pulse simultaneously with a radio message so that receivers can measure the time difference on arrival (TDOA) as described in Cricket [24]. When the transducers are face to face, our implementation can achieve up to 12m range with less than 5cm error. Comparable implementations were able to achieve proportionally similar results of 3-5m range with 1-2cm accuracy [13, 21, 25]. The differences in magnitude are due in part to our design decision to reduce the center frequency of the transducer from the standard 40kHz to just above audible range at 25kHz, which increases both maximum range and error.

Ultrasound transducers are highly directional, and small variations from a direct face to face orientation can have large effects on error and connectivity. Two solutions have been proposed to use ultrasound in multi-hop networks: aligning multiple transducers outward in a radial fashion [21] or by using a metal cone to spread and collect the acoustic energy uniformly in the plane of the other sensor nodes [25]. We implemented the latter solution as shown in Figure 1. In this configuration, our nodes achieve about 5m range and 90% of the errors are within 6.5cm. A comparable implementation achieved about 3m range [25].

All deployments used the Ad-hoc Positioning System's (APS) DV-distance algorithm [16], which is representative of a large class of distributed localization algorithms that use shortest-path [11, 26, 27] or bounding-box [28, 29] approximations. APS uses a distance vector algorithm to approximate the shortest path distance through the multihop network from each node to each of the anchor nodes. Each shortest path distance approximates the true distance to the anchor, reducing the multi-hop localization problem to a single-hop localization problem with a more complex range estimate. The approximate distance to each anchor is then used with the anchor node positions to triangulate the position of each node using linear least-squares.

APS has been shown to yield comparable results to the other distributed localization algorithms [18] and, intuitively, all of these algorithms suffer from the same two sources of error: they will underestimate distances due to long radio hops and will overestimate distances in sparse networks.

In our implementation, the APS algorithm runs in three fully decentralized phases. When the anchor nodes are given their positions, they trigger a *ranging phase* in which all nodes estimate the distance to each of their direct ranging neighbors. The anchors then initiate a *shortest path phase*, in which anchors initiate a tree broadcast, allowing each node to determine its shortest path to each anchor in a distance vector manner. When all broadcasts are complete, each node estimates its position in the *localization phase*. After the anchor nodes were given their positions, the entire process was automated with no human intervention or central computer and completed in less than five minutes for each deployment. All ranging estimates, shortest paths and estimated locations were stored in RAM on the nodes and




Fig. 2. Data Collection Topology. This specially generated topology with 25 nodes measures 300 different distances with at least 1 distance every .025m between 0.4m and 5.2m.

were collected by an automated script after each run.

In all deployments the nodes were placed in a random grid formation, which is like a grid with random noise added to the X and Y coordinates of each grid location. A random grid prevents artifacts of the strict regularity of a grid or of the possible network partitions in a completely random distribution, neither of which would be representative of a canonical deployment. We delegated the four nodes nearest to the corners to be the anchor nodes because keeping all nodes within the convex hull defined by the anchors has been shown to be optimal [2].

The deployment process was non-trivial, especially for RSS localization, and addressed issues of noise characterization, triggering global phase transitions in the network, avoiding collisions during the ranging phase, and reducing the number of retransmissions in the shortest path phase. We also developed several techniques to obtain better results than those presented here. However, neither the implementation issues we faced nor the techniques we developed to increase accuracy is the contribution of this paper. Rather, we focus on identifying the key factors that must be addressed to obtain simulations results that closely model real world deployments.

4. SIMULATION METHODOLOGY

We simulated both the ultrasound and the RSS deployments described in Section 3 and compared the simulation results with the observed deployment results. We use two different techniques for simulation: the traditional technique based on parametric models and a new, more accurate technique that we designed based on statistical sampling. By simultaneously using different simulation techniques, one for rang-

	No Noise	Gaussian Noise	Sampled Noise
Unit Disk Connectivity	D/N	D/G	D/S
Sampled Connectivity	S/N	S/G	S/S

(a) The six kinds of simulation combine three models of noise with two models of connectivity. Each combination is used to simulate the three actual deployments.




Fig. 3. Experimental Results. Each graph in (b), (c) and (d) compares the results of a real-world deployment with each of the six kinds of simulation in Table (a). The box indicates median error; the errors bars indicate upper and lower error quartiles.

ing noise and one for ranging connectivity, we have six different combinations of simulation techniques labeled in Table 3(a). Traditional simulation is used to generate Gaussian noise and Unit Disk connectivity while statistical sampling is used to generate what we call Sampled Noise and Sampled Connectivity. The notation *C/N* stands for the particular *connectivity* and *noise* combination of a simulation. For example, D/G refers to the simulation with Unit Disk connectivity and Gaussian noise. Experiments D/N and S/N in the first column use simulated connectivity but not simulated ranging noise. In this section, we describe the two different simulation techniques used; the results of the simulation experiments will be discussed in the next sections.

In traditional simulation, data is generated from a parametric function. Thus, Gaussian noise is generated for experiments D/G and S/G with the function $\mathcal{N}(d_{ij}, \sigma)$ and Unit Disk connectivity is generated for experiments D/N, D/G, and D/S using the inequality $d_{ij} \leq d_{max}$. For traditional simulation to be meaningful, the model parameters d_{max} and σ should be estimated from empirical ranging data. The typical data collection technique for ranging is to place a transmitter and receiver at several known distances and measure the response [14, 15, 21], although this technique doesn't account for several sources of noise such as node variability. Following the commonly used methodology, in our simulations we used parameters $d_{max} = 20m$, $\sigma = 2m$ for RSS and $d_{max} = 5m$, $\sigma = 6.5cm$ for ultrasound.

We developed an alternative simulation technique based on *statistical sampling* where we generate data for simula-

tion by randomly drawing measurements from an empirical data set. Define the distribution $M(\delta, \epsilon)$ to be the empirical distribution of all observed ranging estimates for distances in the interval $[\delta - \epsilon, \delta + \epsilon]$. We generate a ranging estimate \hat{d}_{ij} for simulation by using the error of a random sample from $M(d_{ij}, \epsilon)$. For example, if \hat{d} is the empirical estimate selected from $M(d_{ij}, \epsilon)$, then

$$\hat{d}_{ij} = d_{ij} + (\hat{d} - \dot{d}_a) \quad (2)$$

where \dot{d}_a is the actual distance at which \hat{d} was measured. Because $\hat{d} \sim M(d_{ij}, \epsilon)$, the simulation is using empirical distributions for signal noise and connectivity as long as $M(d_{ij}, \epsilon)$ accurately represents ranging characteristics at d_{ij} .

The set $M(\delta, \epsilon)$ can include *ranging failures*, which are instances of when a pair of nodes failed to obtain a ranging estimate. Ranging failures are necessary to correctly model ranging connectivity. To generate Sampled Noise alone in experiments D/S and S/S, however, ranging failures are not included in the set. To generate Sampled Connectivity alone in experiments S/N, S/G, and S/S, ranging failures are included, and we define $c_{ij} = \text{true}$ if and only if the sampled ranging estimate \hat{d}_{ij} is not a ranging failure.

The challenge in using this sampling technique is to collect ranging error and connectivity data with a high enough resolution so that small values of ϵ can be used. For example, if we want to use $\epsilon = 2.5cm$ and ultrasound ranging has a maximum range of 10m, we must take empirical ultrasound measurements at 400 different distances. Instead

of measuring each distance with a single pair of nodes, all measurements can be taken at once with $\sqrt{400} = 20$ nodes in a topology where each pair of nodes measures a different distance. By adding a few additional nodes, we can get multiple pairs at each distance. We generated such topologies using rejection sampling, i.e., we generated thousands of topologies until one of them exhibited the desired properties. For example, we used the topology in Figure 2, which required 25 nodes to obtain 2.5cm resolution over 5m, to characterize ultrasound. The topology we used for RSS required 30 nodes to obtain 30cm resolution over 30m.

All nodes are placed at random orientations in this topology and each node transmits 10 times in turn while all other nodes receive. To remove the bias of each distance being measured by only two pairs of nodes (the reciprocal pairs A/B and B/A), we repeated this procedure five times with different mappings of nodes to the topology locations. These mappings were generated using rejection sampling to ensure that the same distances were not always measured by the same pairs. The procedure generated 100 total measurements at each distance with 10 different transmitter/receiver pairs. Therefore, with $\epsilon = 0.05m$ (two inches) the set $M(\delta, \epsilon)$ is likely to include 400 empirical measurements.

Unlike the conventional pairwise technique described above, the empirical measurements in $M(\delta, \epsilon)$ are taken with dozens of transmitter/receiver pairs, capturing a broad spectrum of node, antenna, and orientation variability. Furthermore, the measurements are taken over several different paths through the environment, capturing variability due to dips, bumps, rocks or other environmental factors. Finally, this technique captures connectivity characteristics by fixing the number of transmissions and measuring the number of readings at each distance. In contrast, the conventional pairwise technique described above requires the experimenter to take readings at every possible distance, burying the degradation of ranging connectivity with distance.

The rejection sampling algorithms required on average twelve hours to compute the topology and node mappings. Each data collection process required approximately 6 hours to complete, with the bulk of the time needed for data collection and to precisely measure out the special topology.

5. EXPERIMENTAL AND SIMULATION RESULTS

The ultrasound deployment was repeated 7 times and yielded a median error of 0.78m. The RSS deployments were repeated 10 times each and yielded median errors of 4.3m and 13.4m error for the 49 and 25 node deployments, respectively. Each of the three deployments was simulated with the six simulation combinations shown in Table 3(a), and each simulated experiment was repeated 100 times. Figure 3 compares the median error of the real-world deployments to the median errors of the corresponding simulations. Recall the notation C/N stands for the particular *connectivity* and *noise* combination of a simulation. Also, $D/*$ refers to all simulations with Unit Disk connectivity and $*/G$ refers to all simulations with Gaussian noise.

We can identify the source of error in each deployment by examining which subset of simulations accurately predicts the observed error in each deployment. The 49 node RSS deployment is well predicted by both the $*/G$ and $*/S$ simulations but not the $*/N$ simulations. This trend indicates that noise is the dominant cause of the localization error in this deployment. In contrast, the 49 node ultrasound deployment is well predicted by the $S/*$ simulations but not the $D/*$ simulations. This indicates that the ultrasound connectivity is different than the Unit Disk model, and these deviations dominate noise as the source of error in this deployment. The 25 node RSS deployment shows a similar trend; the $S/*$ simulations predict observed error better than the $D/*$ simulations, but no connectivity/noise combination correctly predicts all the error in this deployment. This indicates that ranging characteristics besides noise and connectivity are causing localization error. The following four sections provide a deeper analysis of these trends.

6. SUFFICIENCY OF NOISY DISK AT HIGH DENSITY

The three empirical deployments fall into two distinct groups: high ranging density and low ranging density, where ranging density is defined by the average degree in the graph defined by successful ranging estimates. A recent study shows that localization is quantitatively different in high ranging density networks than low ranging density networks, and the transition occurs at an average degree of about 7.5 [18]. As density drops below 7.5, localization accuracy increases quickly, but it stabilizes at densities higher than 7.5. According to this criteria, the 49 node RSS deployment has a high ranging density with average degree of 9 while the 49 node ultrasound and the 25 node RSS deployments have low ranging density with average degrees of 6 and 3 respectively.

The 49 node RSS deployment and its corresponding simulations in Figure 3(d) show that the Noisy Disk model is a sufficient model of ranging for deployments with high ranging density. The $*/N$ simulations do not accurately predict observed error, indicating that noise is the dominant cause of localization in this deployment, and the $*/G$ and $*/S$ simulations behave similarly, indicating that RSS noise is sufficiently close to a normal distribution. Furthermore, the $D/*$ and $S/*$ simulations perform similarly, indicating that any differences between Unit Disk and Sampled Connectivity do not influence the localization error. This makes sense at high densities where, even with slightly different types of connectivity, the network should maintain an average degree greater than 7.5, which means the error will be stable.

7. A TRANSITION REGION IN CONNECTIVITY

Our 49 node ultrasound deployment had an average node spacing of 2.2m. With a nominal maximum range of 5m, the Unit Disk model of ultrasound would predict this deployment to have an average degree of 14, which is well




Fig. 4. Ultrasound Transition Region. The gray scale indicates the loss rate (or the level of connectivity); the size of the box indicates the fraction of nodes at that distance with that level of connectivity. The entire range of ultrasound is a transition region; it exhibits neither bimodal nor disk-like connectivity.

above the threshold for a high density deployment. However, an average degree of only 6 was actually observed during deployment, and accordingly, the localization error was 5.7 times worse than predicted by the Noisy Disk model. A comparison between the D/* and S/* simulations indicates that a difference between the Unit Disk and Sampled Connectivity accounts for most of the error in this deployment. Unlike the 49 node RSS deployment where noise was the dominant source of localization error, noise has very little effect in this deployment.

Figure 4 illustrates empirical ultrasound connectivity characteristics, showing the fraction of pairs at each distance that exhibit each of ten levels of connectivity. This figure illustrates what is commonly known as a *transition region*: distances at which some nodes have 100% connectivity while others have 0% connectivity. The unit disk model assumes that nearby nodes have 100% connectivity, far nodes have 0% connectivity, and that the transition region in between is very small. Recent studies have shown that, with low-power radios, the transition region can extend over as much as 50% of the useful radio range, violating the Unit Disk model and introducing problems for networking algorithms that assume disk-like connectivity [30]. Figure 4 shows that ultrasound connectivity is even worse: the transition region extends over the entire range, and there is no distance that clearly defines the difference between connected nodes and unconnected nodes.

The transition region seen in ultrasonic connectivity violates several assumptions made by various localization algorithms about disk-like connectivity. For example, some algorithms assume that all non-connected pairs are farther than some distance d_{max} [31]. However, it is clear from Figure 4 that no such distance exists. Other algorithms assume that all connected pairs will be closer than some dis-

tance d_{max} [32, 33]. While this is true for some value of d_{max} , any such value must be very large relative to the average ranging distance. In our deployments, the ultrasound hardware measured distances more than 50% greater than the nominal maximum range, and other empirical studies have indicated similar findings for radio connectivity [27].

Although APS does not make strict assumptions about disk-like connectivity, it is still greatly affected by the transition region because given a certain maximum range for a ranging technique, a large transition region yields fewer total ranging estimates than the Unit Disk model would predict. This is important to capture in simulation because it affects all localization algorithms by fundamentally reducing the number of constraints on node locations. Modeling this effect would be similar to reducing the maximum range d_{max} , which has been shown to have profound impact on localization [18]. One difference is in the resulting spatial distribution of neighbors: a transition region would result in some far neighbors and some close neighbors, while a small value of d_{max} would result in all neighbors being very close.

8. THE IMPACT OF NON-GAUSSIAN NOISE

While the S/G simulation gets to within 80% of the observed ultrasound error, it is no closer than the simulation S/N, which uses no noise at all. This indicates that the magnitude of ultrasound noise is so small that a Gaussian model of it does not significantly effect localization error. However, S/S arrives to within 94% of empirical error, indicating that a difference between Gaussian noise and Sampled Noise is significantly affecting ultrasonic localization error. While the impact of non-Gaussian noise on localization error is small compared to the effect of non-disk like connectivity, simulations S/G and S/S indicate that it can increase localization error by at least 16%.

The normality plot in Figure 5, in which deviations of data points from the line indicate deviations from the Normal distribution, indicate that ultrasonic ranging generates a heavy-tailed distribution of noise. In other words, it underestimates and overestimates distances more than the Gaussian distribution would predict. This can be detrimental to localization algorithms: the APS algorithm, for example, as the tail with underestimated distances becomes heavier, the shortest-path distances become shorter even if the tail with overestimated distances also becomes heavier. This is because the shortest-path algorithm will selectively ignore paths with many overestimates and will choose the paths with the most underestimates. A similar argument holds for all algorithms that use shortest-path, hop-count [11, 26, 27], or bounding-box [28, 29] techniques.

Similar to the importance of noise itself, the impact of non-Gaussian noise on APS is highly dependent on node degree. A heavy-tailed noise distribution will always serve to shorten shortest-path distance estimates. However, in sparse networks where the shortest-path distances are overestimates due to the zig-zag effect, heavy tailed noise may




Fig. 5. Normality Plot of Ultrasound Noise. The special Y-axis of a normality plot causes normally distributed data to fall in a line. Deviations of data from the line indicate heavy tails in ultrasound noise.

actually decrease shortest path error. In dense networks where the shortest paths are relatively straight, heavy-tailed noise is more likely to increase shortest path error. Furthermore, in dense networks, the shortest path algorithm can choose between many alternative paths, so a smaller number of noisy outliers is necessary to have an impact.

9. BEYOND NOISE AND CONNECTIVITY

Halving the density of the 49 node RSS deployment to 25 nodes creates a low-density RSS deployment that reveals several important insights about RSS localization at low densities. Unlike the ultrasound deployment in which sampled ultrasound connectivity increased localization error by a factor of 4.5, sampled RSS connectivity only increases error by a factor of 2.5. This is likely due to the difference in the respective sizes of the transition region: the transition region for low-power radios is known to be at most 50% of the range, whereas Figure 4 shows that the transition region for ultrasound extends over the entire range. The remaining prediction gap indicates that ranging characteristics beyond noise and connectivity are affecting localization. One such factor may be the non-uniformity of radios or antenna, which would be expected to influence RSS more than ultrasound. Non-uniformity of nodes would be expected to have an effect on connectivity at low densities because some nodes would have very high degree while others would have very low degree, effectively creating partitions in the network. This theory is discussed in Section 10.

A solution to the large prediction gap and the high error observed in this deployment might be to simply increase the transmission power in the network until the average degree of the network is above 7.5. Once the network has high ranging density, a unit disk model should accurately predict localization error. Doing this, however, actually increased error because increasing the transmission power

also increases RSS noise. Increasing the density may close the prediction gap, but it does not necessarily reduce error. The 25 node deployment will always have higher error than the 49 node RSS deployment because RSS ranging is noisier at long distances.

10. STATISTICAL SAMPLING REVISITED

Figure 3 indicates that statistical sampling yields a smaller gap than Noisy Disk at low densities and a similar gap at higher densities. One way to improve simulation predictions is to develop improved ranging models, perhaps borrowing a better model of connectivity from the wireless networking community [34] and extending the Gaussian noise component to include heavy tails for ultrasound. An alternative way is to replace models altogether with statistical sampling. Each approach has its advantages and disadvantages.

Parametric models like Noisy Disk identify a small set of ranging characteristics that affect localization. This provides useful insight into ranging characteristics and the parametric form of the model can be useful in theoretical analysis. One problem with parametric models is that they need to be reevaluated and redeveloped for every new noise characteristic. This is a tedious process requiring data collection and careful analysis followed by a model verification process that may require real localization deployments. Another problem is that empirical ranging characteristics like those shown in Figures 5 and 4 can be too complex to capture in parametric form without some simplification.

Statistical sampling solves both of these problems: new models do not need to be created for new empirical distributions and complex ranging characteristics can be easily captured. However, statistical sampling does not reveal insights about the data nor does it provide a mathematical form that can be used for theoretical analysis.

In practice, parametric modeling and statistical sampling carry similar costs. Both require vast data collection although statistical sampling requires a slightly more careful process. Instead of estimating model parameters from the data, statistical sampling requires the user to generate data subsets $M(\delta, \epsilon)$. During simulation, both methods require a random number to be generated.

The process of creating data subsets $M(\delta, \epsilon)$ is a form of data modeling in the sense that it requires the user to identify which subsets are important, and this method can be extended to model noise characteristics besides noise and connectivity. For example, variations between radios and antennas can be modelled by characterizing the transmitter and receiver type of each radio during data collection. During simulation, each radio could be randomly assigned transmitter and receiver parameters T and R and data could be pooled and drawn from subsets $M(\delta, \epsilon, T, R)$. As long as the parameters T and R are assigned according to the true distribution of radios, this should more accurately model non-uniformity of nodes than using $M(\delta, \epsilon)$.

11. DISCUSSION

This study suggests a top-down approach to evaluating models by comparing each model's predictions with empirical observations of localization deployments. This is in contrast with the commonly used bottom-up approach for deriving models by analyzing raw empirical data [35]. A bottom-up approach is useful for identifying and characterizing the few most important features of empirical data and building them into a model. A top-down approach can evaluate whether the model is a *sufficient* representation of those features, and whether that set of features is sufficient to represent the empirical data.

Our study finds that the Noisy Disk model predicts deployment error fairly well in situations when density is high enough that noise is a dominant factor. However, as density decreases the difference between Unit Disk and empirical connectivity begins to dominate the effects of noise. For RSS ranging, noise and connectivity alone do not sufficiently represent all empirical ranging characteristics; other effects like variations among radios may also effect RSS noise and/or connectivity. These findings provide insight into the empirical nature of ranging characteristics and how they impact localization, suggesting directions for the future design of both ranging models and localization algorithms.

Many multi-hop localization algorithms have yielded extremely accurate results in simulation but there has been a general feeling in the community that obtaining these results is much harder in real-world deployments. To some extent, this study explains why this is true and presents statistical sampling techniques that bring the challenges of the real world into simulation.

12. REFERENCES

- [1] I. Ziskind and M. Wax, "Maximum likelihood localization of multiple sources by alternating projection," in *IEEE Transactions on Signal Processing*, October 1988, vol. 36, pp. 1553–1560.
- [2] A. Savvides, W. Garber, S. Adlakha, R. Moses, and M.B. Srivastava, "On the error characteristics of multihop node localization in ad-hoc sensor networks," in *IPSN*, 2003.
- [3] C Chang and A. Sahai, "Estimation bounds for localization," in *IEEE SECON*, October 2004.
- [4] D. Moore, J. Leonard, D. Rus, and S. Teller, "Robust distributed network localization with noisy range measurements," in *Sensys*, 2004, pp. 50–61, ACM Press.
- [5] Parham Aarabi, "Localization-based sensor validation using the Kullback-Leibler divergence," *IEEE Transactions on Systems, Man and Cybernetics, Part B*, vol. 34, no. 2, pp. 1007–1016, April 2004.
- [6] "TPS: A time-based positioning scheme for outdoor wireless sensor networks," in *IEEE INFOCOM 2004*, 2004.
- [7] Paramvir Bahl and Venkata N. Padmanabhan, "Enhancements to the RADAR user location and tracking system," Tech. Rep. MSR-TR-2000-12, Microsoft Research, February 2000.
- [8] Leon Evers, Stefan Dulman, and Paul Havinga, "A distributed precision based localization algorithm for ad-hoc networks," in *Pervasive Computing: Second International Conference, PERVASIVE 2004*, 2004, pp. 269–286.
- [9] Neal Patwari and Alfred Hero, "Using proximity and quantized rss for sensor localization in wireless networks," in *2nd International ACM Workshop on Wireless Sensor Networks and Applications (WSNA)*, 2003.
- [10] Anton Schwaighofer, Marian Grigoras, Volker Tresp, and Clemens Hoffmann, "Gpps: A gaussian process positioning system for cellular networks," in *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.
- [11] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz, "Localization from mere connectivity," in *MobiHoc*, June 2003.
- [12] Xiang Ji and Hongyuan Zha, "Sensor positioning in wireless ad-hoc sensor networks with multidimensional scaling," in *IEEE INFOCOM 2004*, 2004.
- [13] A. Savvides, H. Park, and M.B. Srivastava, "The bits and fbps of the n-hop multilateration primitive for node localization problems," in *First ACM International Workshop on Sensor Networks and Applications*, 2002.
- [14] M. L. Sichitiu, V. Ramadurai, and P. Peddabachagari, "Simple algorithm for outdoor localization of wireless sensor networks with inaccurate range measurements," in *International Conference on Wireless Networks 2003*, 2003, pp. 300–305.
- [15] N. Patwari, A. Hero, M. Perkins, N. Correal, and R. O'Dey, "Relative location estimation in wireless sensor networks," *IEEE Transactions on Signal Processing, Special Issue on Signal Processing in Networks*, vol. 51, no. 8, pp. 2137–2148, August 2003.
- [16] D. Niculescu and B. Nath, "Ad Hoc Positioning System (APS)," in *IEEE GLOBECOM*, 2001, pp. 2926–2931.
- [17] C. Savarese, J. M. Rabaey, and J. Beutel, "Locationing in distributed ad-hoc wireless sensor networks," in *ICASSP 2001*, 2001.
- [18] K. Langendoen and N. Raijers, "Distributed localization in wireless sensor networks: a quantitative comparison," *Computer Networks*, vol. 43, no. 4, pp. 499–518, November 2003.
- [19] J. Hightower and G. Borriello, "Location systems for ubiquitous computing," *IEEE Computer*, vol. 34, no. 8, pp. 57–66, August 2001.
- [20] J. Hightower, R. Want, and G. Borriello, "SpotON: An Indoor 3D Location Sensing Technology Based on RF Signal Strength," Tech. Rep. UW CSE 2000-02-02, University of Washington, Feb 2000.
- [21] A. Savvides, C. C. Han, and M. B. Srivastava, "Dynamic fine-grained localization in ad-hoc networks of sensors," in *MobiCom*, 2001, pp. 166–179.
- [22] K. Whitehouse, "The design of calamari: an ad-hoc localization system for sensor networks," M.S. thesis, University of California at Berkeley, 2002.
- [23] K. Lorincz and M. Welsh, "A robust, decentralized approach to RF-based location tracking," Tech. Rep. TR-19-04, Harvard University, 2004.
- [24] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *MobiCom*, 2000.
- [25] L.E. Navarro-Serment, C.J.J. Paredis, and P.K. Khosla, "A beacon system for the localization of distributed robotic teams," in *The International Conference on Field and Service Robotics*, Pittsburgh, PA, August 1999.
- [26] C. Savarese, "Robust positioning algorithms for distributed ad-hoc wireless sensor networks," M.S. thesis, University of California at Berkeley, 2002.
- [27] R. Stoleru and J. A. Stankovic, "Probability grid: A location estimation scheme for wireless sensor networks," in *SECON*, 2004.
- [28] S. Simic, "A distributed algorithm for localization in random wireless networks," unpublished manuscript, 2002.
- [29] A. Savvides, H. Park, and M. B. Srivastava, "The bits and fbps of the n-hop multilateration primitive for node localization problems," in *First ACM International Workshop on Sensor Networks and Applications*, September 2002.
- [30] Jerry Zhao and Ramesh Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *The First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.
- [31] P. Biswas and Y. Ye, "Semidefinite programming for ad hoc wireless sensor network localization," in *IPSN*, 2004, pp. 46–54, ACM Press.
- [32] J. Blumenthal, F. Reichenbach, M. Handy, and D. Timmermann, "Optimal adjustment of the coarse grained localization-algorithm for wireless sensor networks," in *WPNC 2004*, March 2004.
- [33] L. Doherty, K. Pister, and L. El Ghaoui, "Convex position estimation in wireless sensor networks," in *INFOCOM 2001*, 2001.
- [34] M. Zuniga and B. Krishnamachari, "Analyzing the transitional region in low power wireless links," in *SECON*, October 2004.
- [35] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic, "Impact of radio irregularity on wireless sensor networks," in *Mobisys*, 2004.