Localization with Snap-Inducing Shaped Residuals (SISR): Coping with Errors in Measurement

H.T. Kung, Chit-Kwan Lin, Tsung-Han Lin, Dario Vlah
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
{htk, cklin, thlin, dario}@eecs.harvard.edu

ABSTRACT

We consider the problem of localizing wireless nodes in an outdoor, open-space environment, using ad-hoc radio ranging measurements, e.g., 802.11. We cast these ranging measurements as a set of distance constraints, thus forming an over-determined system of equations suitable for non-linear least squares optimization. However, ranging measurements are often subject to errors, induced by multipath signals and variations in path loss, unreliable hardware or antenna connectors, or imperfection in measurement models. Such potentially large, non-Gaussian errors in the measurement data ultimately produce inaccurate localization solutions. We propose a new error-tolerant localization method, called *snap-inducing shaped residuals* (SISR), to identify automatically "bad nodes" and "bad links" arising from these errors, so that they receive less weight in the localization process. In particular, SISR snaps "good nodes" to their accurate locations and gives less emphasis to other nodes. While the mathematical techniques used by SISR are similar to robust statistics, SISR's exploitation of the *snap-in* effect in localization appears to be novel. We provide analysis on the principle of SISR, illustrate errors in real-world measurements, and demonstrate a working SISR implementation in field experiments on a testbed of 37 wireless nodes, as well as show the superior performance of SISR in simulation with a larger number of nodes.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*

General Terms

Algorithms, Experimentation, Measurement, Performance, Theory

Keywords

Localization, Least Squares Optimization, Wireless Ranging, 802.11, Robust Statistics

MobiCom'09, September 20–25, 2009, Beijing, China.

1. INTRODUCTION

Power-based localization, which relies on range estimation via received signal strength of RF signals, is convenient in ad hoc networks of wireless nodes. First, ad-hoc Wi-Fi-based localization is infrastructure-free. It has advantages over infrastructure-based systems such as GPS, RFID and cellular telephone localization, when these infrastructures are unavailable, provide insufficient accuracy, or suffer from slow response times. Second, Wi-Fi hardware is already a standard part of many mobile electronic devices. Since localization using Wi-Fi ranging can utilize any Wi-Fi transmissions (e.g., even IBSS beacons), in a sense, location information can be derived for free. Often, stable Wi-Fi ranging measurements can be quickly obtained from a modest number of packets, allowing for faster localization (under a second). This is in contrast to GPS devices which may suffer from much longer bootstrap delays, while waiting to acquire satellite locks.

In this paper, we consider the problem of using 802.11 radio ranging to localize ad-hoc nodes in an outdoor, open-space environment. The solution is straightforward when the terrain is flat and nodes can all hear each other well. However, this may not be the case in real-word applications. For example, a node in a pit, with a faulty antenna, or at the edge of radio range to other nodes may cause large errors in ranging measurements due to a weak signal from the node or received by the node. In addition, heterogeneous radio equipment may also introduce inconsistency in measurement errors. Such errors can create asymmetric ranging measurements between nodes and, often, are non-Gaussian in nature. Just a few such instances can drastically affect the proper use of a measurement model and accuracy of a localization solution.

We propose a technique to cope with the presence of a relatively small number of possibly large, non-Gaussian errors in ranging measurements. The method, "least squares with snap-inducing shaped residuals," or *snap-inducing shaped residuals* (SISR) for short, emphasizes those computed node locations which match well with ranging measurements, while de-emphasizing others. That is, localization based on SISR will favor those computed locations which represent good matches for a majority of measurement data, even if these locations may mean a large deviation from a minority.

2. POTENTIAL APPLICATIONS

We envision that outdoor ad-hoc localization could be well-suited for a host of applications. There are at least five broad application areas where outdoor ad-hoc localization can be naturally applied: coordination, monitoring, tracking, auditing, and wireless routing. We consider these in the context of outdoor areas that are large enough such that localization is useful, but too small for accurate GPS resolution; an example would be an area the size of a football field.

First, outdoor ad-hoc localization could be useful in task coordination. For example, security personnel assigned to an outdoor concert arena could use each other's location information to coordinate the appropriate coverage of the area. Similarly, robots used for exploration of urban areas, where tall buildings make the receipt of GPS signals difficult, could be placed in a team such that tasks are coordinated based on each team member's location.

Second, location information can be integral to outdoor monitoring applications. For example, an elder or child safety monitoring system deployed in a backyard or park could use ad-hoc localization to quickly pinpoint subjects in an emergency. Third, localization can be used in tracking. For example, a directional speaker can send beamformed music to a homeowner walking in her backyard by tracking her position via a Wi-Fi transmitter she wears. In a human speaker identification application, tracking the speaker's location allows the system to extract audio features of high SNR [13]. Monitoring and tracking can also encompass surveillance applications, where there are no expectations that the target be cooperative. We posit that SISR could be used to locate intruders via radio tomography, since human bodies significantly reflect radio waves [24]. Consider nodes, deployed in an outdoor field, that are constantly engaged in SISR localization; in the absence of a human body, the localization solution for these nodes are relatively stable. When an intruder walks across the field, the localization solution of nodes near the person will change significantly since errors will be introduced in their radio ranging measurements. With sufficient node density, it could be possible to track the intruder's path through the deployment.

Fourth, we consider auditing applications, which are a natural extension of monitoring applications. An example would be in work force management, where employees working outdoors would be required to wear radio transceivers so that managers could audit their locations at any moment. An alternative scenario would be in red/blue team training exercises, whether in military applications or in more sportive pursuits such as football, where the location of team members over the course of the exercise is key to useful post-game analysis.

Finally, accurate outdoor localization could be useful in many mobile ad-hoc network (MANET) contexts. MANET routing protocols that rely on signal strength or packet delivery rates can suffer from route flapping due to transient fades in signal or traffic loads. In an outdoor environment, physical node location provides a more stable estimate of link quality and can be used in geometry-aware or geographic routing protocols. Localization information can also be used to perform self-clustering to uncover cluster structures in an ad-hoc network, enabling cluster-driven protocols [5, 12]. In future work, we aim to apply our localization method to ad-hoc computing clusters, where multiple directional antennas or optical transceivers can be dynamically steered to maximize parallel throughput, when given accurate localization solutions.

3. RELATED WORK

There is a large body of literature on the problem of extracting geometry information from a network of wireless nodes, e.g., self-localization in wireless sensor networks. Existing solutions fall into two major categories: range-free and range-based localization. Typically, range-free localization methods [16, 9, 21, 20] utilize node connectivity and hop-count information along with geometric constraints to determine node locations. On the other hand, aiming at providing more accurate localization results, range-based localization methods [19, 18, 2] acquire pairwise distance measurements between wireless nodes and use this information to derive a localization solution. Typical range estimation (ranging) techniques include measuring time-of-arrival (TOA) or received signal strength indications (RSSIs).

A major issue for range-based localization schemes is error in distance (ranging) measurements. Whitehouse et al. [23] have indicated that the errors of RSSI-based RF ranging do not follow a Gaussian distribution, implying that conventional least squares optimization schemes are particularly susceptible to such errors. Without properly taking this into account, localization results can be drastically altered or biased by only a few bad measurements.

Error mitigation in range-based wireless localization has recently received some attention in the literature, with several groups proposing iterative and incremental distributed localization algorithms. Liu et al. [14] use an explicit error management approach to prevent error propagation during incremental localization, and propose a modified least-squares objective function that includes a perturbation term such that the difference in measurement data can be minimized in an average sense, resulting in less sensitivity to errors. Other research efforts have taken a different tack, focusing on identifying and exploiting rigid network topologies to defend against flip and discontinuous flex ambiguities in a localization solution [7, 15]. In this paper, we assume a dense network with a large number constraints, where flip and flexing ambiguities are less likely to happen.

Like SISR, other methods such as semidefinite programming (SDP) [2] and multidimensional scaling (MDS) [21] have taken a centralized approach to the localization problem. Centralized approaches can be adopted in ad-hoc networks with an additional step of collecting data to a master node, but, in comparison with distributed methods, have advantages such as fewer information exchanges between individual nodes and more efficient computation. Reducing inter-node communication messages is especially significant for a wireless ad-hoc network, since wireless medium is shared. SDP maps the localization problem to a convex optimization problem and is able to locate nodes in the presence of errors, but does not consider discounting the outliers. Similarly, MDS-based techniques are robust against some forms of measurement errors. By replacing all-pairs ranging measurements with a multi-hop shortest path distance, some of overestimated distances are excluded or reduced in the localization process. However, this strategy does not perform well for non-convex topologies, where ends of a curved topology can be deformed due to overestimation by the shortest path distance. We will revisit this issue later in the paper.

The SISR estimator is similar to those found in "robust statistics" [8, 10], the study of outlier rejection problems and techniques. These methods have been used in image processing and pattern recognition, for such applications as object velocity estimation from a sequence of images [4], anisotropic diffusion [3], and bilateral filtering [6]. Ultimately, robust estimation techniques share the same goal as outlier detection methods [14]. They both aim at reducing the effect of bad measurements on estimates, but take different approaches. Robust estimation starts from extracting coherent good measurements and disregarding the outliers, while outlier detection methods try to identify and remove individual outliers first. Note that robust estimation can also identify outliers by their relatively large residuals. However, robust estimation has the advantage that finding a sufficient number of good measurements is usually easier than picking out every individual outlier [11].

Although similar, the SISR estimator differs from those previously proposed in robust statistics in some important details of how the residual function is shaped (see Section 5). These differences make SISR localization results more accurate for snapped-in nodes and also make the convergence of the method more robust. To our knowledge, SISR appears to be the first work on applying robust

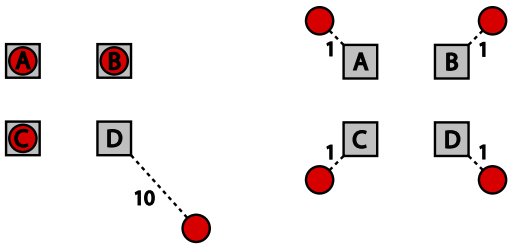

(a) Uneven (b) Even

Figure 1: Two possible localization solutions of nodes A, B, C, and D, where D has large measurement errors. Squares indicate the ground truth location; circles, the computed localization solution. (a) Solutions for A, B, C are very accurate, but that for D is very inaccurate (magnitude "10"), since D's measurement error is not amortized over the other nodes. (b) The measurement error from D is amortized over A, B, and C, causing them to be localized with decreased accuracy (magnitude "1").

statistics to localization. Specifically, we use "snap-in" of good nodes as a localization objective. It appears that our use of the notion of "snap-in" in the context of robust statistics is novel.

4. MOTIVATION OF THE SISR APPROACH

To illustrate the goal of SISR, consider a four node scenario where nodes A, B and C incur small measurement errors on their links, while node D incurs large errors. Thus, A, B and C are "good nodes", and D is a "bad node." A localization method could lead to one of the following two solutions, as shown in Figure 1:

- **Solution *UNEVEN***, which localizes A, B and C accurately, but misses D by a large margin.
- **Solution *EVEN***, which localizes D with the same order of accuracy as A, B, and C, but achieves this at the expense of accurately locating A, B and C.

We prefer solution $UNEVEN$ to $EVEN$, since D cannot be localized accurately in any case, given that its links have large measurement errors. Moreover, since A, B, and C could potentially be localized with great accuracy, a localization method that returns solution $UNEVEN$ ought to de-emphasize D's measurements in order to avoid polluting the localization solutions of A, B, and C. A conventional least squares method would find solution $EVEN$, since it does not differentiate between good and bad ranging measurements.

Developing a localization method that can accurately locate "good nodes" by discerning "bad nodes," is the motivation behind our work in SISR. An assumption we make throughout is that good nodes will represent a majority of nodes. In Section 5.3, we will show a break-down analysis on the maximum percentage of "bad nodes" allowed. SISR makes a key modification to the conventional least squares method: the residual function is *shaped*. Our resulting *snap-inducing shaped residuals (SISR) method* is capable of differentiating between good and bad nodes automatically and, consequently, is able to find the desired solution $UNEVEN$. Furthermore, by applying the method iteratively, SISR can achieve greater accuracy in localization for good nodes.

The failure mode that we consider in this work—that of the "bad node"—differs from the typical failure mode found in the literature, where individual directional links are assumed to experience ranging errors determined by an error model that is independent and identical across links. In practice, physical effects such as multipath phenomena, shadowing or even faulty hardware can produce gross error that is insufficiently captured by such a single, link-scope error model. In addition, such effects can give rise to correlated ranging errors (e.g., a node that has both faulty transmit and receive antennas) which simple, unidirectional link error models cannot capture. The effects of these types of errors are not yet well-understood and, to our knowledge, no localization methods prior to SISR have been developed to specifically defend against such errors.

5. SNAP-INDUCING SHAPED RESIDUALS (SISR)

In this section we present a new kind of residual function for use in optimization-based localization. We will say that the residual function is *snap-inducing* due to its tendency to preserve residuals smaller than some threshold, while diminishing the effect of the larger ones. We give a theoretical analysis of the SISR estimator and show that the properties of SISR can be explained by the theory of robust statistics.

5.1 The SISR Estimator

The localization problem is typically framed as a non-linear least squares optimization, where ranging measurements are used as constraints, and a best fit is sought to minimize squared residuals. In the context of localization, the residual is defined as the difference between pairwise ranging measurements and the distances estimated by the estimator. Eq (1) shows the mathematical formulation of the residual $r(i,j)$ from a ranging measurement between nodes i and j.

$$r(i,j) \quad = \quad \widehat{d}_{ij} - d_{ij} \tag{1}$$

where $\widehat{d}_{ij}$ is the distance estimated by the least-squares estimator and d_{ij} is the ranging measurement between the two nodes.

Conventional least squares works by minimizing an objective function, which is the sum of squared residuals over all node pairs (i,j):

$$F \quad = \quad \sum_{i,j} r(i,j)^2 \tag{2}$$

The squared residual function is shown in Figure 2. The least-squares estimator is not robust to extremely noisy measurements because the squared residual grows quadratically. To solve the problem as well as produce the snapped-in effect for good nodes, we propose a new *snap-inducing* shaped residual (SISR) estimator, which has a shaping function that deemphasizes the influence of bad nodes while emphasizing the good ones. The function is sketched in Figure 2 and has the following two properties.

1. The shaping function increases with a smaller slope when the residual is large. In particular, the function dampens the impact of residuals larger than a threshold τ. We call this the *wing-shaped section*.
2. The shaping function has a narrow and deep well for residuals close to 0. The potential good measurements can therefore be emphasized by growing the shaped residuals or the "cost function" more rapidly. This is called the *U-shaped section*.

For the solutions $EVEN$ and $UNEVEN$ described in Section 4, Property 1 punishes the former by amplifying the residuals in A,

B and C, while Property 2 forgives the latter by diminishing the effect of the increased residual in D. This means SISR will favor the desired solution $UNEVEN$.

The general form of the SISR shaping function is shown in Eq (3):

$$s(i,j) = \begin{cases} \alpha r(i,j)^2 & \text{if } |r(i,j)| < \tau \\ \ln(|r(i,j)| - u) - v & \text{otherwise} \end{cases} \quad (3)$$

where α, τ, u, and v are parameters to be configured.

Tuning the shape of the SISR function controls its sensitivity to errors. The parameter α is introduced to control the height of the U-shaped section, while τ controls its width. (Note that the "U" portion has its x-axis values in the range [-τ, +τ].) The wing-shaped section is created by taking logarithms of the residuals larger than τ. Note that in robust statistics literature, an almost flat function is often used to produce similar robust behaviors (see, e.g., the Lorentzian estimator in Figure 4c [10]). However, a flat objective function with zero slope would mean an arbitrary number of possible solutions, so it could be difficult to handle numerically. Numerical algorithms such as Levenberg-Marquardt require a non-zero gradient of the objective function, and therefore we choose a slowly increasing log function instead. In addition, we note that the use of such an increasing function rather than a flat one will allow good nodes to have more opportunities in finding better snapped-in locations. Another common requirement for numerical methods is that the objective function has to be continuous and differentiable. To make the SISR function piecewise-continuous and piecewise-differentiable at τ, the other two parameters u and v should have the form shown in Eq (4) and (5), respectively.

$$u = \tau - \frac{1}{2\alpha\tau} \quad (4)$$

$$v = \ln(\frac{1}{2\alpha\tau}) - \alpha\tau^2 \quad (5)$$

In short, the SISR function is controlled by two tunable parameters α and τ. The SISR estimator operates by minimizing the objective function as the sum of the shaped residuals over all node pairs (i,j).

$$F = \sum_{i,j} s(i,j) \quad (6)$$

Note that the function for residuals greater than τ need not be logarithmic; for example, we have found that radical functions (e.g., $x^{1/2}$) exhibit the above desired characteristics as well. The trade-off is in running time performance—flatter wings will generally increase the time a numerical method takes to converge to a solution while providing a larger domain of convergence.

In the SISR estimator setting a very small τ can lead to more accurate localization results, but will also increase the likelihood of falling into an incorrect local minimum. On the other hand, a more permissive τ reduces this likelihood at the expense of localization accuracy. In the extreme when τ approaches to infinity, SISR is reduced to the conventional least-squares estimator. We can naturally exploit this trade-off by performing multiple rounds of SISR where τ is successively reduced in each round and the localization solution of one round is used as the initial approximation for the next (see Section 7.1).

5.2 Effect of Measurement Errors on SISR: A Simple Illustrative Simulation

To provide some quick insight into how SISR manages measurement errors, we consider a basic example where k nodes are randomly deployed on a square field, and $k(k-1)$ pairwise ranging measurements are available. Suppose only one of these ranging measurements has a different magnitude of measurement error. Figure 3 shows, in simulation, how such errors affect conventional least-squares localization and SISR localization.

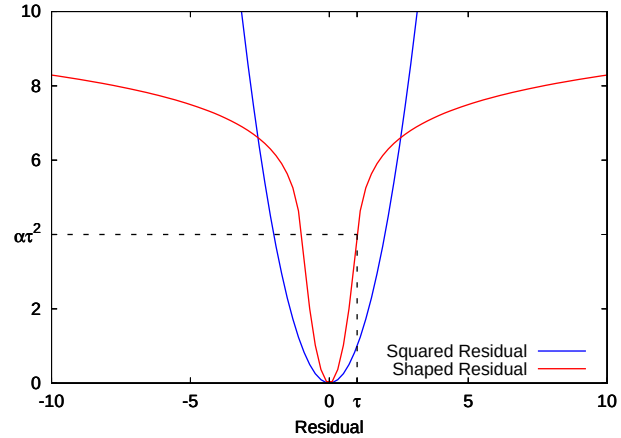

Figure 2: A comparison between the standard squared residual used in conventional least squares and our shaped residual, as in Eq (3) with $\alpha = 4, \tau = 1$ and with u and v as defined in (4) and (5). Note that the shaped residual function is shaped like a "U" near 0 and like "wings" for values larger than τ.

In the simulation, we placed 20 nodes on a 100×100 unit field. The performance of the two localization methods is evaluated using the average pairwise distance error for the $k-2$ nodes that the bad link is not incident on. We can see that the error obtained by least-squares continues to grow with the increasing error in the bad link. This is the standard least-squares behavior whereby large residuals are reduced by increasing any other smaller residuals. In effect, the bad link is causing the computed topology to flex, and so we call this a *flexed localization*. The error bars are not shown for clarity; however, the least-sqares result does have a large variance, because not every randomly chosen bad link has the same ability to flex the rest of the topology.

The SISR error behaves similarly up to a point. Specifically, as long as the distortion in any given residual is smaller than the threshold τ, according to (3), the residual's contribution will be same as least-squares. However, any distortions larger than τ will contribute diminishing residuals, leaving the topology "snapped into" its original shape; we can see this effect in Figure 3 as the SISR error departs from the least-squares error. In fact, as the bad link's error grows, its effect on the rest of the nodes *decreases*—this is because the larger the bad link's residual, the less we can decrease it—and thereby deform the good links—before the net cost starts to increase.

5.3 Analysis of SISR and Breakdown Points

Here, we give a theoretical analysis of the SISR estimator and show how the snap-in effect emerges in SISR. Recall parameter τ is introduced to control the estimator's error sensitivity. τ is the residual value at which the "U" portion of the shaping function $s(r)$ attains its maximum. The shaped residual at τ is denoted as $s(\tau)$. To facilitate analysis, we introduce parameter η, defined as the maximum shaped residual we expect to encounter in a given localization scenario. By denoting the maximum residual as r_{upper}, then $\eta = s(r_{upper})$. Since the wing portion of the SISR shaping function is slowly increasing, we assume that the ratio of η to $s(\tau)$ is a small positive constant value $c > 1$. That is, $\eta = c \cdot s(\tau)$. Generally, we can assume that the value of c is 2 to 3. In our sim-

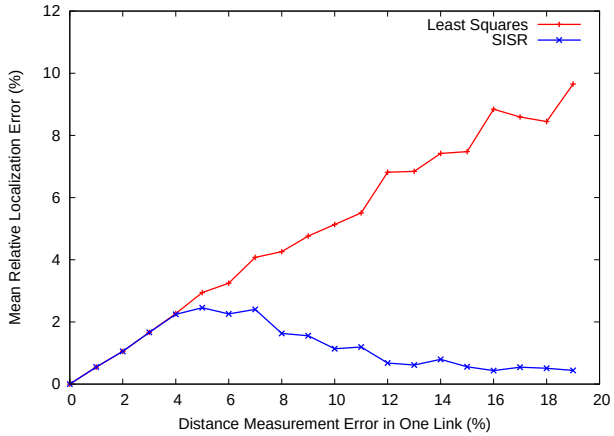

Figure 3: A simple illustrative experiment comparing localization performance of least-squares vs SISR under exactly one erroneous link. Each point is an average of 100 runs. The least-squares mean error grows along with the input error, while the SISR result remains "snapped" into place. (Note that localization error starts descending when measurement error is beyond a certain level, due to the decreasing slope of the "wings" portion of the SISR shaping function.)

ulations, where $\tau = 0.25$ and $\alpha = 80$ and the largest expected residual is 25m, $c = 2.38$. Even if we increase the largest expected residual to 100m, $c = 2.65$.

We use the same scenario as in Section 5.2 where k nodes are deployed on a field. Suppose that the ranging measurements for all the $k(k-1)$ links are accurate, and that all k nodes have been correctly localized without any error. That is, all the $k(k-1)$ links have zero residuals.

Now suppose that a node n_k changes its ranging measurements with the other $k-1$ nodes, with the change in each ranging measurement being no more than r_{upper}. We compare the estimation of the two different estimators: the SISR estimator, and the least-squares estimator. We refer to the localization solution using SISR as *snapped-in localization*, and that using least-squares as *flexed localization*.

(a) **Snapped-in localization.** All k nodes are snapped into their original locations, i.e., there are no changes in their computed locations. This means that distances between node n_k and the other $k-1$ nodes no longer match the corresponding changed ranging measurements. Thus, there is an increase in each of the $(k-1)$ residuals by an amount no greater than r_{upper}. In terms of the shaped residual for each link, each increase is no more than $s(r_{upper})$, i.e., η. Thus, the total increase in shaped residuals, $\Delta r_{snapped}$, must satisfy (7):

$$\Delta r_{snapped} \quad \leq \quad (k-1)\eta \tag{7}$$

(b) **Flexed localization.** In this solution, we assume here that all the k nodes change their locations by at least τ in distance, in order to adjust to the changed ranging measurements involving node n_k. Furthermore, we assume that the location changes are not synchronized in their directions, so each pair of the k nodes will observe a change in their distance by at least τ. Recall that only node n_k has changed its ranging measurements, and the ranging measurements for pairs in all the other nodes remain unchanged. Thus, each of these node pairs incurs an increase in its residual by at least τ. In terms of shaped residuals, the increase is at least $s(\tau)$ or η/c. Since there are $(k-1)(k-1)$ ranging measurements subject to change, the total increase in shaped residuals satisfies (8):

$$\Delta r_{flexed} \geq (k-1)^2 \frac{\eta}{c} \tag{8}$$

Note that the condition for $\Delta r_{flexed} > \Delta r_{snapped}$ is:

$$k > c + 1 \tag{9}$$

Thus, when k is sufficiently large, snapped-in localization will incur a smaller cost than flexed localization. This means that the SISR estimator has the desired property that it will find the snapped-in localization as a local minimum.

Following the same argument, we can also show that the snap-in effect occurs when there are multiple nodes changing their ranging measurements. In the case where p nodes change their ranging measurements rather than only one node n_k, Δr_{flexed} and $\Delta r_{snapped}$ can be shown to be:

$$\Delta r_{snapped} \quad \leq \quad p(k-1)\eta \tag{10}$$

$$\Delta r_{flexed} \quad \geq \quad (k-p)(k-1)\frac{\eta}{c} \tag{11}$$

Suppose that b is the value of p at which Δr_{flexed} equals to $\Delta r_{snapped}$. Then b is the *breakdown point* of SISR. If there are more than b nodes that change their ranging measurements and the changes can be arbitrarily large, then SISR can no longer snap into the correct location. By equating the bounds of (10) and (11), we can solve for a lower bound on the breakdown point:

$$b \geq \frac{k}{1+c}. \tag{12}$$

From (12), we note that, under the model of our analysis, SISR can snap-in to good nodes, as long as the percentage of bad nodes is under $k/(1+c)$. When c is close to 1, i.e., the wing portion of the SISR shaping function is almost flat, SISR can tolerate approximately 50% nodes being bad. Since no estimator can tolerate more than 50% bad measurements, we can conclude that in that case the breakdown point is exactly at 50%. In our formulation with a logarithmic wing portion, c is typically 2 to 3, and the breakdown point is 25% to 33%, or higher. Table 1 gives lower bounds on breakdown points for different c values.

Table 1: Lower Bounds for Breakdown Points of SISR

c	1	1.5	2	2.5	3
Breakdown pt.	50%	40%	33%	29%	25%

Above, we assume that in flexed localization, all nodes flex to minimize the total residuals. This assumption can be relaxed. For example, our argument still holds by requiring only a sub-linear number $f(k)$ of nodes to flex rather than all k nodes, as long as $f(k)$ is a monotonically increasing function, e.g., $f(k) = \log k$ or $\sqrt{k}$. Further, we note that for random node placements, location changes under least squares optimization are indeed generally not synchronized in their directions (see Figure 3), as we assumed in the analysis.

5.4 SISR and Robust Statistics

The SISR estimator can be viewed under the lens of robust statistics, a theoretical framework concerned with the outlier rejection problem in statistical analysis. Outliers can bias the final outcome of an estimator, lead to incorrect conclusions, and are usually hard to filter out—especially when different outliers of different magnitudes are present. Estimators which suffer in the presence of outliers are mean and least squares, in the sense that the estimation can be altered without bound by an extremely noisy outlier. In contrast,

the median estimator is not as susceptible to such polluting data, and is considered a robust estimator.

To assess the robustness of an estimator, an associated *influence function* can be used to characterize the importance of individual data samples [8]. A smaller absolute value in the influence function means the data receives less weight in the estimation. The influence function is proportional to the derivative of the estimator. For example, Figures 4a and 4b show the least-squares estimator and its influence function. The influence function grows linearly with the residual of a sample and, as a result, a single noisy data sample can have extreme effect on the estimator. On the other hand, the *Lorentzian* estimator [8] , shown in Figure 4c, has a *redescending influence function*, shown in Figure 4d. The function approaches 0 for large residuals, and thus data with large residuals have diminishing effects on the estimator. Such estimators are less sensitive to gross errors.

The SISR estimator and its influence function are illustrated in Figures 4e and 4f, respectively. The influence function has a redescending shape similar to that of the Lorentzian, and we know SISR is also a robust estimator. SISR departs from the Lorentzian in that SISR can take advantage of certain special assumptions on the structure of underlying data samples. In the context of localization, SISR works best when the pairwise distance measurements are either close to 100% accurate or far off. To address this scenario, SISR uses a residual shaping function which has a narrower center well and flatter wings at the two ends. With such a shaping function, SISR can ignore the bad data and stick with the good measurements. The flatter wings also give SISR a higher breakdown point than the Lorentzian because of a smaller c value. Our experimental results in Section 8 show that SISR indeed gives better localization estimates than the Lorentzian estimator.

Applying SISR to two-dimensional (2D) node localization problems is particularly attractive because of its snap-in effect. Nodes situated on a 2D space are constrained by geometry and thus have less freedom to move than nodes in a higher dimensional space. In other words, the number of good measurements needed by SISR to produce the "snapped-in" effect in 2D is fewer. This benefit can be seen in the theoretical analysis on SISR breakdown points. As noted earlier, under the model used in our analysis SISR can tolerate as much as 50% of total measurements being in error.

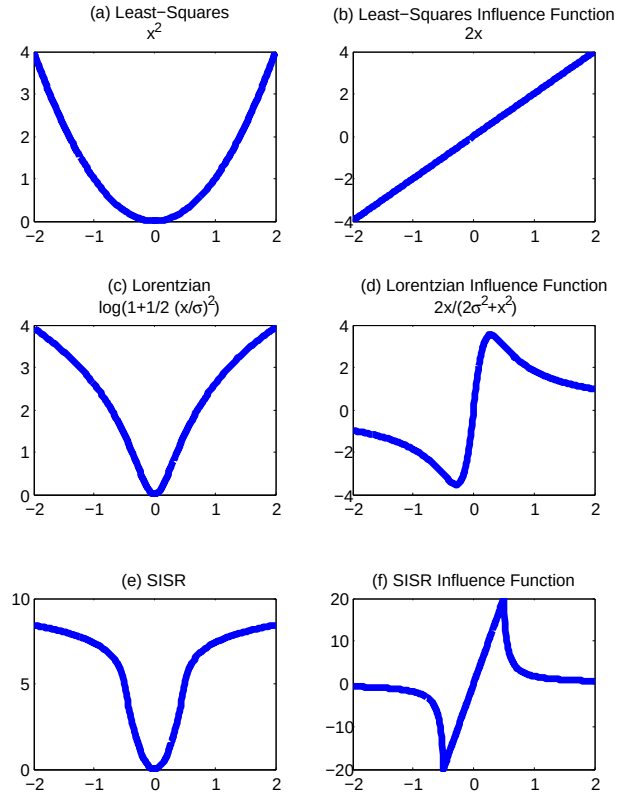

Figure 4: Robust estimators and their influence functions.

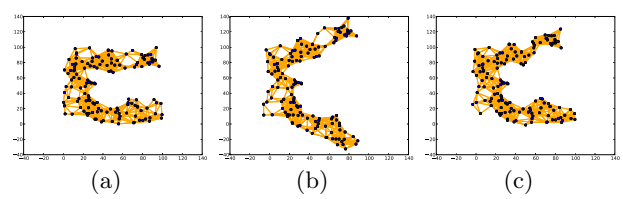

Figure 5: (a) A non-convex "C"-shaped topology with 147 randomly-placed nodes within a 100m × 100m bounding area. Topologies such as these were used as ground truth in simulations described in Section 7. (b) The corresponding MDS-MAP(C,R) solution. Deformation in the ends of the "C" are clearly visible. (c) The corresponding SISR solution, which corrects for the deformation, even when ranging errors are present.

6. EVALUATION METHODOLOGY

We evaluate SISR localization efficacy against that of an optimized variant of multidimensional scaling—MDS-MAP(C,R) [21]. A key advance in MDS-MAP is the replacement of the ranging measurement matrix with a "proximity matrix" P, where p_{ij} is the shortest path distance between nodes i and j. This mitigates (a common) ranging error where distances between nodes are overestimated and, in multi-hop topologies, provides a proxy ranging estimate for node pairs that are not within communication range. MDS-MAP(C,R) further develops this innovation by performing a weighted least squares refinement of a classic MDS-MAP solution, where the weights are inversely proportional to the number of hops in the shortest path. This strategy dampens the effect of long-range estimates, which are more susceptible to errors, but makes the blanket assumption that such estimates are always less trustworthy. In contrast, SISR can discern between good and bad ranging measurements and automatically weigh down those that are bad.

In general, MDS-MAP is an effective method for multi-hop localization. However, for certain non-convex multi-hop topologies the shortest path distance gives rise to errors in localization, which is a known limitation [22, 1]. For example, in the "C" topology in Figure 5a, the true distance between nodes at the ends of the "C" is much smaller than the shortest path distance. On such a topology, MDS-MAP causes the ends of the "C" to repel (Figure 5b), and ranging errors further exacerbate the problem. For this reason, we focus on these difficult non-convex topologies in our simulations.

Even for such topologies, given a reasonably good initial approximation SISR can arrive at a very accurate localization solution for nodes that have good ranging measurements. Since the MDS-MAP(C,R) method does not depend on an initial approximation, its result is a good candidate for the initial approximation to SISR. Therefore, we use it as such in our simulations.

The most natural way to evaluate the performance of these algorithms is through an *absolute localization error metric*, i.e., the discrepancy between a node's computed location and its ground truth location. However, the formulation of SISR we have pre-

sented thus far is an anchor-free, relative localization scheme; it depends only on pairwise distance measurements and does not return its results in the same coordinate system as the ground truth. The core MDS-MAP(C,R) algorithm is also such a scheme.

In order to report absolute localization error, we modified both algorithms slightly to employ the ground truth locations of four anchor nodes. As in [21], the anchors are randomly selected. For MDS-MAP(C,R), the solution returned by classic MDS-MAP is first transformed (via rotation, translation and reflection) into the ground truth coordinate system by fitting the anchors to their ground truth locations. This provides a better-conditioned initial solution to the subsequent refinement stage, where the anchor locations are further used as additional, fixed-constraint equations passed to the weighted least squares process. This forces the least squares operation to respect the ground truth coordinate system. The same technique is employed in providing constraint equations to SISR. Note that the way we use anchor information is different from the method in [21], where the final localization solution is fitted to the ground truth via anchors. Since we postulate the availability of anchor information, it makes sense to take advantage of it during the localization procedure.

In 2D localization, three anchors are sufficient, but their random selection can cause ambiguities during coordinate system fitting (e.g., if the three anchors are almost collinear). In [21], a fourth anchor is used to reduce the ambiguity, but there is no guarantee that this will always work. Thus, a question arises: what criteria describe a good set of anchors? First, they must be well-separated. Second, the angles between any pair must not be near zero. We note that, in its anchor-free, relative localization formulation, SISR can be used to bootstrap an absolute localization by providing hints as to which nodes satisfy these two criteria: (1) SISR's snapped-in effect allows accurately localized nodes to be distinguished by their low residuals, and (2) we can distinguish which among these low-residual nodes are well-separated in the relative localization solution. However, for fairness in comparison with MDS-MAP(C,R), we do not employ this method in our evaluations below.

7. SIMULATION RESULTS

In our first set of simulations, we examine how much SISR can improve upon MDS-MAP(C,R) in non-convex, "C"-shaped topologies, under various error conditions. Given a topology and set of error conditions, we first compute the MDS-MAP(C,R) localization solution. This is then used as an initial approximation for SISR localization. Finally, we compare the accuracy of the solutions from the two methods.

For each trial, we generate a topology where 147 nodes[1] are placed uniformly at random within a "C"-shaped region in a bounding area of 100m $\times$ 100m. An example of such a topology is shown in Figure 5a. Each node has a simulated radio range of 16.5m, which results in a multi-hop topology.

Each trial is also described by two experiment conditions governing error. The first condition is the percentage of "bad nodes" within the population (0%, 10%, 20%, 30%, 40%, 50%). A node is assigned to either a "good" population, in which case it has no ranging error, or a "bad" population, when it experiences ranging error $\epsilon \sim N(\mu_{bad}, 0.2\mu_{bad})$. The second condition is μ_{bad} (5%, 10%, 20%, 30%, 40%, 50%). Note that a "bad node" has ranging error only from itself to other nodes and not vice versa; this results in asymmetric range measurements, a more realistic scenario.

[1]We chose 147 nodes to maintain the same node density in the "C"-shaped region as 196 nodes placed in a regular grid pattern in the same area.

We have 36 combinations across the two experiment conditions, for a total of 36 simulations. Each simulation consisted of 20 trials on distinct topologies. The average median absolute localization errors are shown Figure 6. First, consider the results in Figure 6a, where MDS-MAP(C,R) is shown to be relatively insensitive to all values of ranging error, when the percentage of "bad nodes" is low (<30%), but never gives particularly accurate results. All MDS-MAP(C,R) solutions had at least 2m error, even when there are no "bad nodes" in the population.

In contrast to MDS-MAP(C,R), SISR is capable of providing much more accurate results—less than 2m error—when the percentage of "bad nodes" is less than 30%, as shown in Figure 6b. For these conditions, due to (1) the "snap-in" behavior of SISR, (2) the fact that "good nodes" have no measurement error, and (3) a good initial approximation afforded by MDS-MAP(C,R), the resulting localization solution is not significantly affected by the magnitude of the ranging error. Of course, in a real setting, we cannot expect "good nodes" to be absolutely correct in their range estimates; we will explore one such real setting in Section 8, where we describe our field experiments.

The simulation results above are consistent with our analysis of the SISR estimator's breakdown point in Section 5.3. For $\tau = 0.25$ and $\alpha = 80$, as used in our simulation, the breakdown point for a fully-connected network should be approximately 30% or higher. However, our simulated "C"-topologies were not fully-connected but multihop. Thus, the calculated 30% breakdown point is an upper bound for our simulation scenarios, which agrees with the data presented in Figure 6b.

7.1 Iterative Refinement with SISR

Decreasing τ can improve SISR's accuracy in locating good nodes, provided that good initial approximations are available. This means that we can perform iterative refinement by successively decreasing values of τ. To evaluate the effectiveness of iteratively refining τ, we perform the following simulation: 38 nodes are randomly placed within a "C"-shaped topology within a bounding area of 50m $\times$ 50m (radio range is 16.5m), giving the same node density as in our previous simulations. Bad nodes are assigned ranging measurement error $\epsilon \sim N(0.5, 0.1)$, with a different percentage of bad nodes across three experiments (10%, 30%, 50%). τ is initially set at 5m ($\alpha = 1/\tau^2$, so that the height of the "U"-shaped region does not change). SISR is then run for 10 iterations, with τ decreasing by 25% each iteration. As before, the first iteration is given the result of MDS-MAP(C,R) as its initial approximation. Note that with fewer nodes (as compared to the previous simulations), the optimization problem is actually more difficult, since there are fewer constraints.

Figure 7 shows the result of this simulation. Each data point is the average median localization error in meters over 10 trials. Note that as we decrease τ, the accuracy of the localization improves. For 10% bad nodes, SISR already localizes with high accuracy, even when $\tau = 5$m. The improvement is greatest for 30% bad nodes, where there is room for improvement by SISR and the fraction of bad nodes is under SISR's breakdown point. The improvement is not as marked for 50% bad nodes, as the high number of bad nodes places a fundamental limit on the localization accuracy.

In this simulation, we arbitrarily chose to decrease τ by 25% per iteration for simplicity. In practice, this rate could be tuned by keeping track of the number of low-residual nodes from iteration to iteration (i.e., nodes that snapped-in). In the event that, from one iteration to the next, a large number of these nodes suddenly exhibit higher residuals, it is likely that the reduction rate of τ was too aggressive and that it should be scaled back.

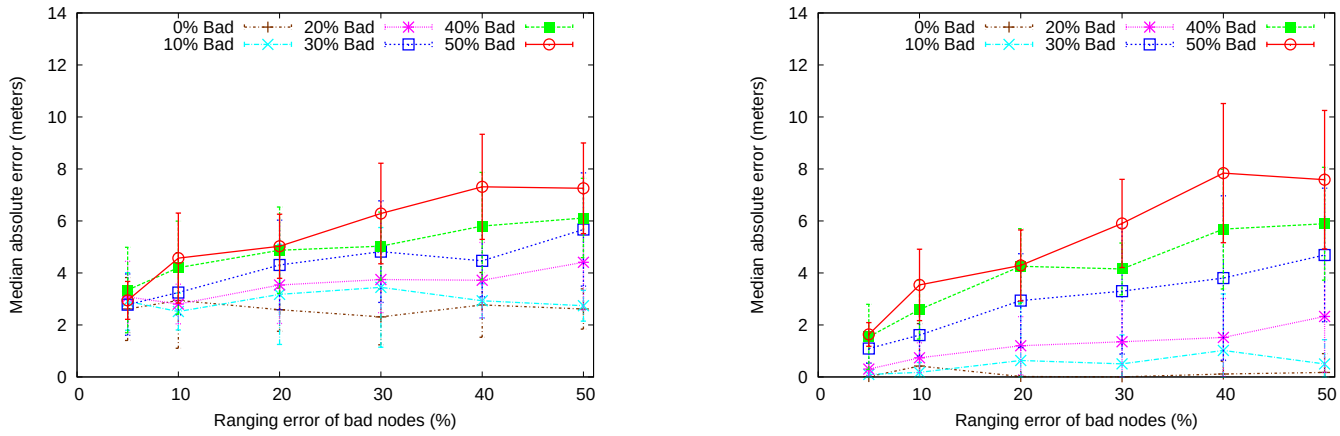

(a) MDS-MAP(C,R) (b) SISR

Figure 6: A comparison of the accuracy of (a) MDS-MAP(C,R) and (b) SISR. Topologies were generated with 147 nodes placed uniformly at random in a "C"-shape within a 100m $\times$ 100m bounding area. For each experiment, a different percentage of nodes were simulated as "bad", i.e., experiencing ranging errors. We vary the magnitude of the ranging error and compare the average median absolute localization error returned by the different methods (20 trials per data point; error bars show one standard deviation). Compared to MDS-MAP(C,R), SISR gives less error in the solution when the percentage of "bad nodes" is $< 30\%$, due to the "snapped-in" effect.

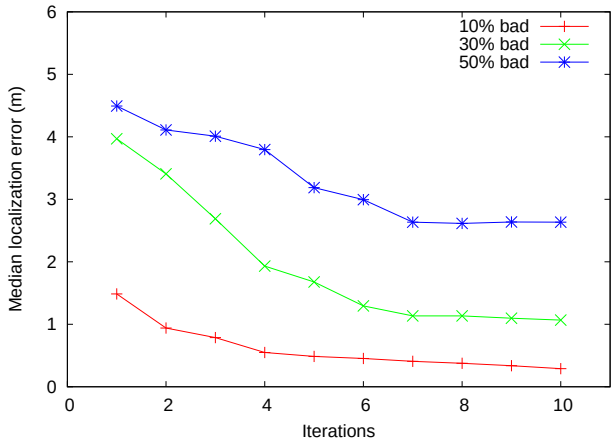

Figure 7: Iterative refinement of localization solutions by gradually decreasing τ gives improved accuracy. Each curve is a separate experimental condition, differing in the percentage of nodes in the population with ranging errors of 50% (10%, 30%, 50% bad nodes).

Similarly, in this simulation, we arbitrarily terminated the τ refinement after 10 iterations. In theory, we could have further reduced τ to produce even better results, since the "good nodes" experience no ranging error. However, in a real-world setting, the accuracy of range measurements on "good nodes" establishes a lower bound on τ (i.e., we cannot expect better performance than our best measurements). Thus, the iterative approach provides a way of tuning τ to the optimal level—when the localization solution no longer improves (as measured by the residuals of the nodes), we know we have reached an optimal value for τ.

8. FIELD EXPERIMENT AND RESULTS

In this section we describe a localization experiment using actual 802.11 radios in an outdoor environment. The goal of this experiment was to validate SISR localization on measurement data collected with real systems and radios as well as study the cause and nature of measurement errors in real-world scenarios (see further discussion in Section 9). We first describe our experimental methodology and then compare the localization results returned by MDS-MAP(C,R), the Lorentzian estimator, and SISR. To achieve a fair comparison between the Lorentzian and SISR estimators, we applied iterative refinement to each. As we will see, SISR effectively mitigates the effect of bad nodes and achieves significant performance gains over MDS-MAP(C,R) and the Lorentzian estimator.

8.1 Experimental Setup

For our experiment site, we chose a flat grass field with no major buildings or other obstacles nearby, and only faint ambient beaconing traffic from some 802.11 access points in the environment. Given the non-negligible effect of objects on link quality [12], we ensured that the area under measurement was clear of personnel and all items other than power cables. We deployed two types of nodes in our outdoor testbed: 21 One Laptop Per Child (OLPC) computers, which contain an x86-compatible AMD processor, and 16 Mobile Internet Devices (MIDs), which are based on the Intel Atom chipset. Both run Linux and have Marvell 802.11b/g radios internally (SDIO-based SD8686 for MID, USB-based 8388 for OLPC). We placed the radios into ad-hoc mode, but were unable to disable 802.11 beaconing and IBSS formation; however, the beacons were relatively infrequent and did not pose significant interference. To obtain ranging measurements, we modified the Linux kernel to provide Received Signal Strength Indicator (RSSI) values as metadata alongside any received UDP packets.

We performed a series of basic one-way signal strength measurement operations between pairs of nodes as follows. One node of the pair would transmit for a relatively long period of time, broadcasting a stream of 4-byte UDP probe packets at the 2Mbps modulation and on 802.11 channel 4. The receiver node would listen for these probe packets, logging each packet's transmitter address and RSSI. Note that it is convenient to perform these basic measurement operations in parallel, by having multiple receivers cap-

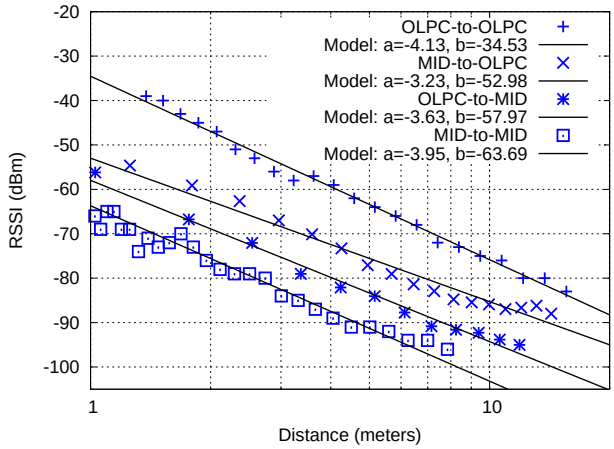

Figure 8: Training data points for the four node type combinations are shown overlaid on the resulting models. The legend contains the model parameters a and b for each combination.

ture a single transmitter's probes. Due to environmental effects, RSSI measurements may be subject to fading. Further, the radio adapters' RSSI reporting mechanism is undocumented, and may report values which deviate from the actual received power. Therefore, in order to obtain a more reliable RSSI estimate, we use a simple windowing heuristic to filter the RSSI data. For node i receiving from node j, we take n total RSSI samples, divide these into w consecutive windows, take the maximum RSSI value in the window and then take the median of these values to be the "representative" RSSI of transmissions from node j. The more samples taken, the more accurate the representative RSSI. From our field data, we have found that the RSSI heuristic converges in as few as 1000 packets ($\sim$ 300ms transmission time), suggesting that the wireless ranging approach we have adopted can be done quickly.

8.2 RSSI-Distance Models

To convert RSSI into range measurements, we chose a path-loss model based on the well-known two-ray propagation [17], and trained it via field measurements. Since all of the links we encountered during the experiment campaign had a much smaller height than length, they operated in the $1/d^4$ regime of two-ray propagation where d denotes the communication distance. Thus, we used the following simple expression for the RSSI-distance model:

$$RSSI(d) = 10\log_{10}(d^a) + b \tag{13}$$

for some constants a and b. We trained the model separately for the four combinations of transmitters and receivers possible with two node types, that is, MID–MID, MID–OLPC, OLPC–MID, and OLPC–OLPC. The height was uniform per node type, so we dropped it from the standard two-ray model and let its effect become subsumed in the parameter b. We performed the training on the same field as the localization measurements, but in different locations; namely, using two nodes, we placed one node (receiver) at 20 points along a straight line, ranging in distance between 0 and 23 meters, while keeping the other node (transmitter) static. A logarithmic spacing was used in placing the 20 points. The resulting measurements and model parameters are shown in Figure 8.

8.3 Pairwise Measurements in the Field

We arranged the 37 testbed nodes (Figure 9) into a "C"-topology, within a 25m $\times$ 25m square region. This area is larger than the

Figure 9: Photo of outdoor testbed with 21 OLPCs (upper left) and 16 MIDs (lower left). These nodes are configured in a "C" topology, on a 25m $\times$ 25m open field (right).

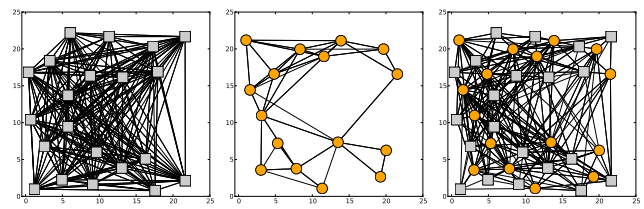

(a) OLPC (b) MID (c) Both

Figure 10: Connectivity maps of the outdoor testbed (a) amongst only OLPCs (squares), (b) amongst only MIDs (circles), and (c) amongst both. Generally, a "C"-topology is exhibited in these maps.

broadcast domain of nodes placed flat on the ground (the diagonal, at $\sim$36m, is well beyond transmission range). With the nodes placed uniformly at random, we expect that each node can only connect to a subset of the other nodes. With the two types of nodes used, there are four possible types of links: OLPC-OLPC, MID-MID, MID-OLPC, and OLPC-MID.

To get a sense of the network connectivity, we plot links between node pairs, according to link type: Figure 10a shows links between OLPCs, 10b shows links between MIDs, and 10c shows links between both (ignoring directionality). MID-MID links are sparser than the others, while OLPCs generally enjoy better connectivity with other OLPCs. This is likely due to OLPC antennas being vertically-oriented and 8cm above the ground, whereas MID antennas are inside the machine enclosure and less than 2.5cm off the ground. We can visually confirm that the network, though not strictly a "C"-topology as depicted in Figure 5a, is indeed "C"-like and not fully-connected.

We measured the signal strength of 1332 ($= 37\times36$) transmitter-receiver combinations (666 pairs, in both directions). We present this data in Figure 11, shown separately for each link type, and together with the corresponding link model.

8.4 Localization Experiment

The pairwise ranging measurements described above were then fed into each of the three localization methods. First, we performed MDS-MAP(C,R), the results of which were then used by the Lorentzian estimator and SISR as initial approximations. For

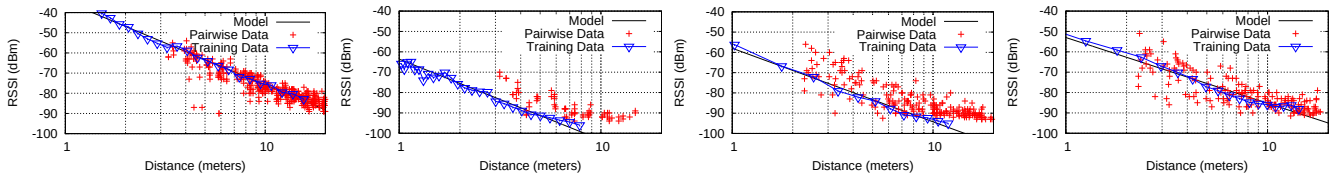

(a) OLPC-OLPC (b) MID-MID (c) OLPC-MID (d) MID-OLPC

Figure 11: Pairwise signal strength measurements from a 37-node 2D localization scenario (Sec 8.3) superimposed on the RSSI-distance model with its training data. There are four separate figures, one for each possible transmitter and receiver type combination.

these latter two methods, we performed iterative refinement to obtain optimal tuning parameters, as described in Section 7.1. With the Lorentzian (recall Figure 4c and 4d), we started with $\sigma = 1.0$ and reduced it by 25% every iteration. With SISR, we initialized $\tau = 5.0$m and similarly reduced it by 25% every iteration. Optimal refinement was achieved after four iterations with the Lorentzian and five with SISR.

Of the nodes in our testbed, we noticed that the MIDs, when communicating amongst themselves, tended to deviate from the MID-MID RSSI-distance model significantly. As discussed in Section 9, this is partly due to MIDs' low-profile antennas sitting just above the ground. This means MID-MID connectivity is more sensitive to antenna orientations and terrain variations. As a result, this subset of MID-MID links introduce a significant amount of errors into the ranging measurements based on the RSSI-distance model, and represents a good test scenario in which to compare the localization performance of SISR to that of MDS-MAP(C,R) and the Lorentzian estimator. Deviations of a trained RSSI-distance model from a specific device and application environment at hand represent a major source of ranging errors. This is an important issue that we will discuss further in Section 9.

Figure 12 shows the ground truth locations of OLPCs (squares) and MIDs (circles) in our field experiment. The lines emanating from the ground truth locations are the absolute localization error vectors. Each one terminates at the localization solution for a node given by each of the methods, respectively: SISR (red), Lorentzian (grey), and MDS-MAP(C,R) (blue). Figure 12 visually verifies that nodes are, on the whole, more accurately localized by SISR. Note that the SISR error vectors are typically shorter for OLPC nodes, but longer for MIDs. This is consistent with our analysis of the behavior of SISR: since the MIDs did not fit the MID-MID RSSI-distance model well, their ranging measurements had large errors and, as a result, SISR dampened their effect on the localization solutions of the other "good" nodes, while allowing the MID solutions to retain high residuals.

Specifically, SISR gives a median absolute localization error of 2.13m, which is lower than that of the Lorentzian (3.12m) and MDS-MAP(C,R) (5.00m). The CDF of absolute localization error across nodes in Figure 13 shows that SISR (red) gives significantly more accurate solutions to a greater fraction of nodes than the other two methods. For example, with SISR, roughly 67% of nodes have 3m error or less. In contrast, only 53% and 33% of nodes achieve the same accuracy under the Lorentzian and MDS-MAP(C,R), respectively. Suppose we examine only the localization performance for OLPC nodes, which incur smaller errors in ranging than MID nodes. Then, SISR outperforms the other two methods by even larger (relative) margins: SISR gives 0.97m median absolute localization error, whereas Lorentzian gives 1.21m, and MDS-MAP (C,R) gives 1.96m. These results are expected, as SISR favors nodes which have better measurements. In fact,

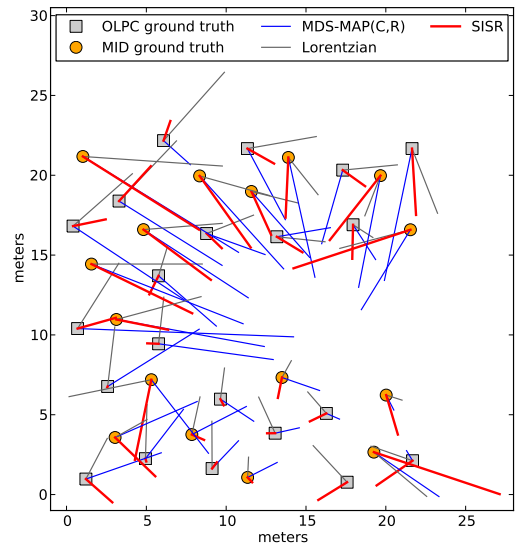

Figure 12: SISR (red), Lorentzian (grey) and MDS-MAP(C,R) (blue) localization error vectors. OLPC (square) and MID (circle) ground truth locations were arranged in a C-like topology. The median absolute localization errors are: SISR = 2.13m, Lorentzian = 3.13m, MDS-MAP(C,R) = 5.00m. SISR outperforms both Lorentzian and MDS-MAP(C,R).

SISR's performance in this case is close to the best possible, given that OLPCs have inherent errors in ranging with a mean magnitude of 1.91m (Figure 11a).

It is interesting to note that the Lorentzian CDF curve eventually crosses the SISR curve. This is also expected behavior—once SISR has snapped in all the "good" nodes, the "bad" ones will remain bad. This results in a steep initial rise in the CDF curve, but a long tail, just as we see in Figure 13. In contrast, the Lorentzian is more permissive of the effect of "bad" nodes on the overall solution, which leads to generally higher error for the "good" nodes (shallowly rising curve), but relatively less error for the "bad" nodes (hence, the cross-over point). This shows that not all robust estimators solve the localization problem equally well and that SISR is particularly suited for the task.

9. PRACTICAL CONSIDERATIONS ON MODELING AND MEASUREMENT

Earlier we saw in Figure 11 that the pairwise measurements deviate in various ways from the training data sets and the derived RSSI-distance models. The reasons for the deviations could be

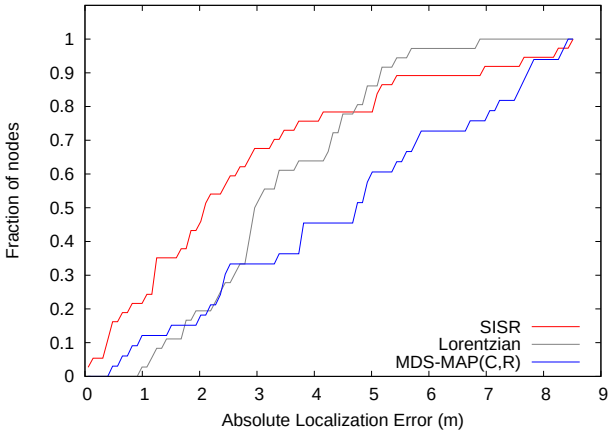

Figure 13: CDFs of the absolute localization error across all nodes as given by SISR (red), the Lorentzian estimator (grey) and MDS-MAP(C,R) (blue).

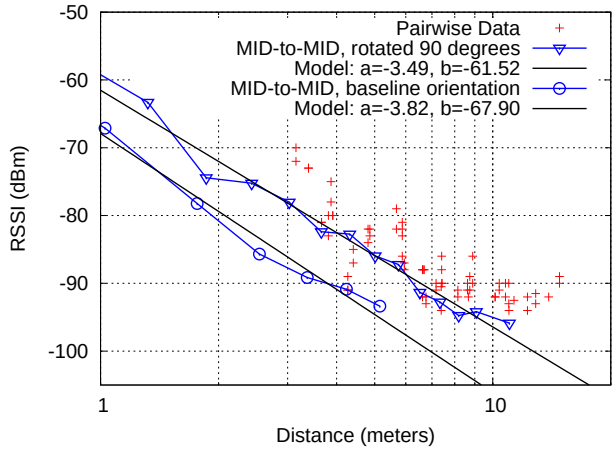

Figure 14: Two MID-MID training data sets taken for the devices oriented as in all the rest of the measurements, and rotated by 90°.

non-flat ground surface, asymmetric antenna patterns, the effect of objects in the measurement area, or manufacturing variation. We will discuss these in turn.

Ground Curvature. The two-ray propagation model assumes a perfectly flat ground surface. Given our nodes' proximity to the ground (especially, MIDs), even small deviations from a flat surface may have significant impact on measurements. To verify this effect, we performed additional training measurements for a MID-MID and OLPC-MID link, along four straight, parallel paths crossing the field at 7.6m intervals; for each of these paths we determined the curvature of the ground by measuring the height of a horizontal laser beam at 16 points along each path. The resulting data sets–omitted here to save space–turn out to exhibit a decrease in signal strength corresponding to "hills" in the line-of-sight signal path, and vice-versa.

The effect of curvature is less severe for the OLPC-MID link, because the height of the hills in the signal paths—up to 7cm, and largely around 2cm—is smaller than the height of the OLPC antennas—8cm. Therefore, higher-mounted antennas are beneficial not only for improved radio range, but also smaller modeling errors.

Asymmetric Antenna Patterns. To gauge the effect of antenna patterns, we repeated one of the MID-MID training runs with both devices rotated by 90° in the ground plane; the resulting data set is shown in Figure 14 together with the original one. We see that antenna orientation has a consistent and significant effect on the strength of a MID-MID link. For nodes deployed in a 2D space, it is unlikely for a node with a fixed antenna orientation to match, that is, have maximal antenna gain to all its surrounding neighbors. This means that it is inevitable that some node pairs exhibit deviations from the RSSI-distance model, no matter how accurate the model is.

Consider our "C"-shaped placements of MID nodes in Figure 10b, where all MID nodes have their antennas point to the same direction such that horizontally neighboring nodes can receive each other's signal well, while vertically neighboring nodes receive each other worse. These configurations correspond to the two orientation scenarios from Figure 14. Roughly speaking, each node in any specific sub-area of the "C"-shape region of the Figure 10b, such as an end or elbow sub-region, will have some distinct number of matching nodes within its radio reach. The total number of these matching node pairs will be substantially less than the number of all node pairs, which is $16 \times 15 = 240$. The remainder of the node pairs are expected to have lowered RSSI values.

Figure 11b indeed exhibits these expected behaviors in the sense that only 70 node pairs have sufficiently large RSSI values which lead to their inclusion in the plot. (Note that some of these points overlap in the plot. This can happen especially with the two directions of the same link.) It is also instructive to note that in Figure 11b almost all node pairs exhibit RSSI values above the model line, unlike other scenarios depicted in Figures 11a, 11c and 11d. This illustrates the fact that models can not be perfect in spite of a large amount of training.

Objects in the Measurement Area. Objects near the signal path may affect the link quality by either blocking a signal, or providing additional signal paths through reflection. For example, it has been noted in the field that it is possible for links to operate only in the presence of a nearby person [12]. While we kept all personnel and unnecessary items away from the experiment area, there were still objects such as power cables, power strips, or AC/DC power adapters needed to operate the testbed which might have affected link characteristics in similar ways, and were not present during training runs.

Manufacturing Variation. Nodes are generally not physically identical; in the case of radios, for example, the connectors and cables between the amplifier and the antenna typically have variation up to some tolerance specified by the manufacturer. Furthermore, variations outside of that range are possible due to defects. For example, one of the nodes in our testbed turned out to have half the average radio range, but only on reception. It is apparent that a lengthy calibration of individual nodes would be needed to account for such variation; in our case, we only performed the model training measurements with a few of the available nodes.

In general, there are various environmental factors which contribute to loss of model accuracy. Based on the preceding discussion it seems it is possible to account for some of the factors by adding detail to the link models, but at the cost of loss of models' generality and lengthier training. Note that even if exhaustive calibration of nodes is possible, it is often impractical, especially in the field. In these situations, it could make more sense to use error-tolerant localization methods, such as SISR, to address the problem.

10. CONCLUSIONS

In this paper, we present SISR, an error-tolerant localization method for wireless nodes that employ RF ranging. Via analysis, simulation and field experiments, we have demonstrated that our SISR localization method can automatically tolerate errors in ranging measurements by de-emphasizing bad data while accentuating more accurate readings. In particular, SISR can accurately localize difficult non-convex topologies, even in the presence of non-Gaussian errors in ranging. Moreover, SISR's tunable τ parameter allows even higher accuracy in localizing good nodes via iterative refinement. The principle and implementation of SISR are applicable to other ranging methods that experience gross error, e.g., acoustic time-difference-of-arrival (TDOA). Due to limited space, we have not presented our TDOA results here.

We should emphasize that, often, we cannot, and do not wish to, train "perfect" RSSI-distance models for specific application scenarios due to the time and effort involved and the desire for more generic models. SISR is an error-tolerant localization method which can cope with these imperfections in model-based ranging.

That SISR can discount "bad" nodes from the localization process suggests that SISR is also robust against malicious nodes that purposefully report erroneous ranging measurements. "Good" nodes remain snapped-in to their locations while the malicious nodes succeed in ruining only their own localization solutions. But, as with similar optimization techniques, SISR can fail when malicious nodes compromise its initial approximation such that it is placed squarely into an incorrect local minimum at the outset; in these situations, methods which randomize initial approximations (e.g., simulated annealing) can help optimizers break out of such local minima.

In the future, we plan to extend SISR to distributed and cluster-aware localization, where local, reliable clusters would be localized first, and later stitched together by using SISR on inter-cluster ranging measurements. SISR also suggests a natural way of pruning bad nodes from the network: by comparing the unshaped (conventional least squares) residual of nodes, outliers can be rejected. Such node discrimination may be critical to the performance of higher-level network applications. Finally, we will consider using SISR in other applications beyond localization, such as those traditionally studied in the literature by robust statistics.

11. ACKNOWLEDGEMENTS

This material is based on research sponsored by Air Force Research Laboratory under agreement numbers FA8750-08-1-0220, FA8750-09-2-0180 and FA8750-08-1-0191. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

12. REFERENCES

[1] J. Bachrach and C. Taylor. Localization in sensor networks. *Handbook of Sensor Networks: Algorithms and Architectures, 1st ed.*, 1, 2005.

[2] P. Biswas, T. C. Liang, C. K. Toh, Y. Ye, and T. C. Wang. Semidefinite programming approaches for sensor network localization with noisy distance measurements. *IEEE Trans. Autom. Sci. Eng.*, 3(4):360–371, 2006.

[3] M. Black, G. Sapiro, D. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Trans. on Image Processing*, 7(3):421–432, Mar 1998.

[4] M. J. Black and P. Anandan. The robust estimation of multiple motions: parametric and piecewise-smooth flow fields. *Comput. Vis. Image Underst.*, 63(1):75–104, 1996.

[5] C.-M. Cheng, H. T. Kung, C.-K. Lin, C.-Y. Su, and D. Vlah. Rainbow: A wireless medium access control using network coding for multi-hop content distribution. *MILCOM*, 2008.

[6] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *SIGGRAPH*, 2002.

[7] D. K. Goldenberg, P. Bihler, R. Y. Yang, M. Cao, J. Fang, S. A. Morse, and B. D. O. Anderson. Localization in sparse networks using sweeps. In *MobiCom*, 2006.

[8] F. Hampel, E. Ronchetti, P. Rousseeuw, and W. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. Wiley New York, 1986.

[9] T. He, C. Huang, B. Blum, J. Stankovic, and T. Abdelzaher. Range-free localization schemes for large scale sensor networks. In *MobiCom*, 2003.

[10] P. Huber. *Robust statistics*. Wiley New York, 1981.

[11] M. Hubert, P. J. Rousseeuw, and S. V. Aelst. High-breakdown robust multivariate methods. *Statistical Science*, 23(1):92–119, 2008.

[12] H. T. Kung, C.-K. Lin, T.-H. Lin, S. J. Tarsa, and D. Vlah. Multi-site data exchange over wireless ad-hoc networks using network coding. In *MILCOM*, 2009.

[13] H. T. Kung, C.-K. Lin, C.-Y. Su, D. Vlah, J. Grieco, M. Huggins, and B. Suter. A computational wireless network backplane: Performance in a distributed speaker identification application. In *MILCOM*, 2008.

[14] J. Liu, Y. Zhang, and F. Zhao. Robust distributed node localization with error management. In *MobiHoc*, 2006.

[15] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. In *SenSys*, 2004.

[16] D. Nicolescu and B. Nath. Ad-Hoc Positioning Systems (APS). In *GLOBECOM*, 2001.

[17] J. D. Parsons. *The Mobile Radio Propagation Channel*. John Wiley & Sons, LTD., second edition, 2000.

[18] C. Savarese, J. M. Rabaey, and K. Langendoen. Robust positioning algorithms for distributed ad-hoc wireless sensor networks. In *USENIX*, 2002.

[19] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *MobiCom*, 2001.

[20] A. Savvides, H. Park, and M. B. Srivastava. The n-hop multilateration primitive for node localization problems. *Mob. Netw. Appl.*, 8(4):443–451, 2003.

[21] Y. Shang, W. Ruml, Y. Zhang, and M. Fromherz. Localization from connectivity in sensor networks. *IEEE Trans. Parallel Distrib. Syst.*, 15(11):961–974, 2004.

[22] C. Whitehouse. The design of calamari: an ad-hoc localization system for sensor networks. Master's thesis, UC Berkeley, 2002.

[23] 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.

[24] M. Youssef, M. Mah, and A. Agrawala. Challenges: device-free passive localization for wireless environments. In *MobiCom*, 2007.