

SPASELOC: AN ADAPTIVE SUBPROBLEM ALGORITHM FOR SCALABLE WIRELESS SENSOR NETWORK LOCALIZATION*

MICHAEL W. CARTER[†], HOLLY H. JIN[‡], MICHAEL A. SAUNDERS[§], AND YINYU YE[§]

Abstract. An adaptive rule-based algorithm, SpaseLoc, is described to solve localization problems for ad hoc wireless sensor networks. A large problem is solved as a sequence of very small subproblems, each of which is solved by semidefinite programming relaxation of a geometric optimization model. The subproblems are generated according to a set of sensor/anchor selection rules. Computational results compared with existing approaches show that the SpaseLoc algorithm scales well and provides excellent localization accuracy.

Key words. sensor localization, semidefinite programming, large-scale optimization

AMS subject classifications. 49M37, 65K05, 90C30

DOI. 10.1137/040621600

1. Introduction. Ad hoc wireless sensor networks may contain hundreds or even tens of thousands of inexpensive devices (sensors) that can communicate with their neighbors within a limited radio range. By relaying information to each other, they can transmit signals to a command post anywhere within the network. They have many practical uses in areas such as military applications [15], environment or industrial control and monitoring [7, 9], wildlife monitoring [24], and security monitoring [15]. For example, Southern California Edison’s Nuclear Generating Station in San Onofre, CA, has deployed wireless mesh networked sensors from Dust Networks, Inc. to obtain real-time trend data [9]. These data are used to predict which motors are about to fail, so they could be preemptively rebuilt or replaced during scheduled maintenance periods. The use of a wireless sensor network saves the station money and avoids potential machine shutdown. Implementation of a sensor localization algorithm would provide a service that eliminates the need to record every sensor’s location and its associated ID number in the network.

Wireless sensor networks are potentially important enablers for many other advanced applications. A huge variety of applications lie ahead. By 2008, there could be 100 million wireless sensors in use, up from about 200,000 in 2005, according to the market-research company Harbor Research. The worldwide market for wireless sensors, it says, will grow from \$100 million in 2005 to more than \$1 billion by 2009 [18]. This is motivating great effort in academia and industry to explore effective ways to build sensor networks with feature-rich services [12].

One of the important inputs these services build upon is the exact locations of all sensors in the network. The need for *sensor localization* arises because accurate

*Received by the editors December 27, 2004; accepted for publication (in revised form) February 27, 2006; published electronically December 5, 2006.

<http://www.siam.org/journals/siopt/17-4/62160.html>

[†]Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON, Canada M5S 3G8 (carter@mie.utoronto.ca).

[‡]Department of Management Science and Engineering, Stanford University, Stanford, CA 94305-4026, and Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON, Canada M5S 3G8 (hollyjin@stanford.edu). The author was partially supported by Robert Bosch Corporation.

[§]Department of Management Science and Engineering, Stanford University, Stanford, CA 94305-4026 (saunders@stanford.edu, yinyu-ye@stanford.edu).




FIG. 1.1. Indexing of sensors and anchors.

locations are known for only some of the sensors (which are called *anchors*). If the networks are to achieve their purpose, the locations of the remaining sensors must be determined. One approach to localizing these sensors with unknown locations is to use known anchor locations and distance measurements that neighboring sensors and anchors obtain among themselves. The mathematical problem is to *estimate* all sensors' locations using a sparse data matrix of noisy distance measurements. This leads to a large, nonconvex, constrained optimization problem. Large networks may contain many thousands of sensors, whose locations should be determined accurately and quickly.

1.1. Problem definition. Sensor localization for ad hoc wireless sensor networks aims to find the locations of all sensors in the network, given pairwise distance measurements among some of the sensors and known locations of some of the sensors. The sensors with known locations are called *anchors*. From now on, *sensor* generally means *unlocalized sensor*, excluding anchors. A *node* is any sensor or anchor.

We use a constrained optimization approach to estimate the sensors' locations. The following input, output, and objectives are considered.

Input

Total points: n , the total number of nodes in the network.

Unknown points: s sensors, whose locations $x_i \in \mathcal{R}^2$, $i = 1, \dots, s$, are to be determined. (We assume the points are on a plane here, but the approach is extended to three dimensions in Jin's thesis [14].)

Known points: m anchors, whose locations $a_k \in \mathcal{R}^2$, $k = s + 1, \dots, n$, are known. (Note that we put anchors at the end of the total points' list without loss of generality, and that $n = s + m$. Index k is specific for indexing anchors. Refer to Figure 1.1 for node indexing.)

Known distance measurements: The nonzero elements of a sparse matrix \hat{d} containing the readings of certain ranging devices for estimating the distance between two points. \hat{d}_{ij} is the distance measurement between two sensors x_i and x_j ($i < j \leq s$), and \hat{d}_{ik} is the distance measurement between some sensor x_i and anchor a_k ($i \leq s < k$). The distance measurements are constant data and generally have errors.

Output

Locations: Estimated locations x_i for s sensors.

Objectives

Accuracy: Minimal errors in the estimated sensor locations.

Speed: Fast enough for real-time applications (e.g., networks with moving sensors).

Scalability: Suitable for large-scale deployment (with tens of thousands of nodes).

1.2. Notation. The Euclidean distance between two vectors v and w is defined to be $\|v - w\|$, where $\|\cdot\|$ always means the 2-norm. Nodes are said to be *connected* if the associated measurements \hat{d}_{ij} or \hat{d}_{ik} exist. The remaining elements of \hat{d} are zero.

If a measurement does exist between node i and j but it is zero (i and j are at the same spot), we do not set \widehat{d}_{ij} to zero: we set it to machine precision ϵ instead to distinguish from the case of $\widehat{d}_{ij} = 0$ when two nodes' distance is beyond the sensor device's measuring range.

1.3. Related research work. Sensor localization in ad hoc wireless networks has been a booming research area recently. Hightower and Boriello [12] give an extensive review of the area and available methods. There are many ways to solve the localization problem [6, 8, 10, 13, 17, 19, 20, 21, 22], with two main ones based on triangulation and optimization.

Triangulation methods estimate node locations based on distance measurements between neighboring nodes, and some algorithms use iterative steps to localize all sensors.

Early work using optimization techniques is reported by Doherty, El Ghaoui, and Pister [8]. Ideally the Euclidean distance between neighboring nodes should be fitted in some near-equality sense to the distance measurements:

$$(1.1) \quad \|x_i - x_j\| \approx \widehat{d}_{ij} \quad \text{and} \quad \|x_i - a_k\| \approx \widehat{d}_{ik}.$$

Doherty, El Ghaoui, and Pister formulate a convex optimization model by treating the constraints as $\|x_i - x_j\| \leq \widehat{d}_{ij}$ and $\|x_i - a_k\| \leq \widehat{d}_{ik}$, and by including certain other convex constraints. This formulation takes advantage of available optimization algorithms, including those for convex optimization. However, the method needs sufficient anchors to be on the boundary of the localization area for it to work effectively.

Biswas and Ye [2] work with the near-equality constraints (1.1), and more importantly introduced a semidefinite programming (SDP) relaxation method in order to retain the benefits of convex optimization. They report that their method yields more accuracy than the approach in [8].

The SDP relaxation approach can solve small problems effectively. The paper [2] reports a few seconds of laptop execution time for a 50-node localization problem. However, the number of constraints in the SDP model is $O(n^2)$, where n is the number of nodes in the network. Even a few-hundred-node problem leads to excessive memory and computation time by available SDP solvers such as DSDP (Benson, Ye, and Zhang [1]) and SeDuMi (Sturm [23]). These solvers are effective for SDP problems with dimension and number of constraints up to a few thousand.

Tseng [25] has presented a second-order cone programming (SOCP) relaxation model that permits solution for problem sizes up to a few thousand using available SOCP solvers. However, the additional relaxation of the original model usually generates larger error rates, and the run-times are high. The author reports CPU times of 330 seconds for 1000 nodes and 3 hours for 2000 nodes using SeDuMi 1.05 [23] and MATLAB 6.1 on a Linux PC.

Biswas and Ye [3] propose a decomposition scheme to overcome the scalability issue with SDP solvers. The anchors in the network are first partitioned into many clusters according to their physical locations, and sensors are assigned to these clusters if they have a direct connection to one of the anchors. Each cluster formulates a subproblem, and the subproblems are solved *independently* on each cluster using the SDP relaxation of [2]. The paper reports results for randomly generated sensor networks of 4000 sensors partitioned into 100 clusters strictly according to their geographic locations. Sensors with distance connections to more than one cluster are included in multiple clusters. The final estimation of their locations is determined

by the cluster that gives the least estimated errors. An execution time of about 4 minutes on a 1.2GHz Pentium laptop is reported for a problem of this size. Thus, the decomposition approach makes large-scale sensor network localization possible on a single processor. The further advantage is that multiple CPUs can be used in a natural way.

1.4. SpaseLoc. A basic tool that we have developed during this research is a rule-based iterative algorithm named SpaseLoc (**sub**problem **al**gorithm for **sen**sor **loc**alization). It is effective for networks involving tens of thousands of sensors and beyond, using a single CPU.

To solve a large localization problem (defined as the *full_problem*), SpaseLoc proceeds iteratively by estimating only a portion of the total sensors' locations at each iteration. Some anchors and sensors are chosen according to a set of rules. They form a sensor localization *subproblem* that can be treated similarly to the basic SDP formulation of Biswas and Ye [2]. The solution from the subproblem is fed back to the *full_problem* and the algorithm iterates again until all sensors are localized.

Computational results show that SpaseLoc can solve small or large problems with excellent accuracy and scalability. It is capable of localizing 4000 nodes with great accuracy in under 20 seconds, and 10000 nodes in under a minute on a 2.4 GHz laptop.

2. The subproblem SDP model. This section reviews the quadratic programming formulation of the sensor localization problem and the SDP relaxation model of Biswas and Ye [2] that the SpaseLoc subproblem is based on. Error analysis is also reviewed here as a reference for later sections.

2.1. Euclidean distance model. Consider a network of sensors and anchors labeled as in Figure 1.1. For any point in the network, there could be three types of distance measurements. Since we generally do not need the distance information between two anchor points, we exclude this type of measurement from now on.

The other types of distance measurements are the two we need for the localization model. First is the distance measurement between two sensors (i and j) with unknown locations; second is the distance measurement between a sensor (i) and an anchor (k) with known location. Corresponding to these two types of distances, we define sets N_1 , \bar{N}_1 , N_2 , and \bar{N}_2 as follows:

- N_1 includes pairwise sensors (i, j) if $i < j$ and there exists a distance measurement \hat{d}_{ij} :

$$N_1 = \{(i, j) \text{ with known } \hat{d}_{ij} \text{ and } i < j\}.$$
- \bar{N}_1 includes pairwise sensors (i, j) with unknown measurement \hat{d}_{ij} and $i < j$:

$$\bar{N}_1 = \{(i, j) \text{ with unknown } \hat{d}_{ij} \text{ and } i < j\}.$$
- N_2 includes pairs of sensor i and anchor k if there exists a measurement \hat{d}_{ik} :

$$N_2 = \{(i, k) \text{ with known } \hat{d}_{ik}\}.$$
- \bar{N}_2 includes pairs of sensor i and anchor k with unknown measurement \hat{d}_{ik} :

$$\bar{N}_2 = \{(i, k) \text{ with unknown } \hat{d}_{ik}\}.$$

The full set of nodes and pairwise distance measurements form a graph $G = \{V, E\}$, where $V = \{1, 2, \dots, s, s + 1, \dots, n\}$ and $E = N_1 \cup N_2$.

Introduce α_{ij} to be the difference between the measured squared distance $(\hat{d}_{ij})^2$ and the squared Euclidean distance $\|x_i - x_j\|^2$ from sensor i to sensor j . Also, let α_{ik} be the difference between the measured squared distance $(\hat{d}_{ik})^2$ and the squared Euclidean distance $\|x_i - a_k\|^2$ from sensor i to anchor k . Intuitively, we seek a solution for which the magnitude of these differences is small.

Lower bounds r_{ij} or r_{ik} are imposed if $(i, j) \in \overline{N}_1$ or if $(i, k) \in \overline{N}_2$. Typically each r_{ij} or r_{ik} value is the radio range (also known as *radius*) within which the associated sensors can detect each other.

Biswas and Ye [2] formulate the sensor localization problem as minimizing the ℓ_1 -norm of the squared-distance errors α_{ij} and α_{ik} subject to mixed equality and inequality constraints:

$$\begin{aligned}
 & \underset{x_i, x_j, \alpha_{ij}, \alpha_{ik}}{\text{minimize}} && \sum_{(i,j) \in N_1} |\alpha_{ij}| + \sum_{(i,k) \in N_2} |\alpha_{ik}| \\
 & \text{subject to} && \|x_i - x_j\|^2 - \alpha_{ij} = (\widehat{d}_{ij})^2 \quad \forall (i, j) \in N_1, \\
 & && \|x_i - a_k\|^2 - \alpha_{ik} = (\widehat{d}_{ik})^2 \quad \forall (i, k) \in N_2, \\
 (2.1) \quad & && \|x_i - x_j\|^2 \geq r_{ij}^2 \quad \forall (i, j) \in \overline{N}_1, \\
 & && \|x_i - a_k\|^2 \geq r_{ik}^2 \quad \forall (i, k) \in \overline{N}_2, \\
 & && x_i, x_j \in \mathcal{R}^2, \quad \alpha_{ij}, \alpha_{ik} \in \mathcal{R}, \\
 & && i, j = 1, \dots, s, \quad k = s + 1, \dots, n.
 \end{aligned}$$

The above model is a nonconvex constrained optimization problem. As yet there is no effective solution method. In the following subsections, we review Biswas and Ye’s [2] relaxation method for solving this problem approximately.

2.2. The Euclidean distance model in matrix form. The distance model (2.1) is reformulated into (2.2) (refer to Biswas and Ye [2]) by introducing matrix variables as follows:

$$\begin{aligned}
 & \underset{}{\text{minimize}} && \sum_{(i,j) \in N_1} (\alpha_{ij}^+ + \alpha_{ij}^-) + \sum_{(i,k) \in N_2} (\alpha_{ik}^+ + \alpha_{ik}^-) \\
 & \text{subject to} && e_{ij}^T Y e_{ij} - \alpha_{ij}^+ + \alpha_{ij}^- = (\widehat{d}_{ij})^2 \quad \forall (i, j) \in N_1, \\
 & && \begin{pmatrix} e_i \\ -a_k \end{pmatrix}^T \begin{pmatrix} Y & X^T \\ X & I \end{pmatrix} \begin{pmatrix} e_i \\ -a_k \end{pmatrix} - \alpha_{ik}^+ + \alpha_{ik}^- = (\widehat{d}_{ik})^2 \quad \forall (i, k) \in N_2, \\
 (2.2) \quad & && e_{ij}^T Y e_{ij} \geq r_{ij}^2 \quad \forall (i, j) \in \overline{N}_1, \\
 & && \begin{pmatrix} e_i \\ -a_k \end{pmatrix}^T \begin{pmatrix} Y & X^T \\ X & I \end{pmatrix} \begin{pmatrix} e_i \\ -a_k \end{pmatrix} \geq r_{ik}^2 \quad \forall (i, k) \in \overline{N}_2, \\
 & && Y = X^T X, \\
 & && \alpha_{ij}^+, \alpha_{ij}^-, \alpha_{ik}^+, \alpha_{ik}^- \geq 0, \\
 & && i, j = 1, \dots, s, \quad k = s + 1, \dots, n,
 \end{aligned}$$

where

- $X = (x_1 \ x_2 \ \dots \ x_s)$ is a $2 \times s$ matrix to be determined;
- e_{ij} is a zero column vector except for 1 in location i and -1 in location j , so that

$$\|x_i - x_j\|^2 = e_{ij}^T X^T X e_{ij};$$

- e_i is a zero column vector except for 1 in position i , so that

$$\|x_i - a_k\|^2 = \begin{pmatrix} e_i \\ -a_k \end{pmatrix}^T (X \quad I)^T (X \quad I) \begin{pmatrix} e_i \\ -a_k \end{pmatrix};$$

- Y is defined to be $X^T X$;
- The substitutions $\alpha_{ij} = \alpha_{ij}^+ - \alpha_{ij}^-$ and $\alpha_{ik} = \alpha_{ik}^+ - \alpha_{ik}^-$ are made to deal with $|\alpha_{ij}|$ and $|\alpha_{ik}|$ in the normal way.

2.3. The SDP relaxation model. The approach of Biswas and Ye [2] is to relax the constraint $Y = X^T X$ to be $Y \succeq X^T X$, for which an equivalent matrix inequality is (Boyd et al. [5])

$$(2.3) \quad Z_I \equiv \begin{pmatrix} Y & X^T \\ X & I \end{pmatrix} \succeq 0.$$

With the definitions

$$A_I = \begin{pmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \quad b_I = \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix},$$

where $\mathbf{0}$ in A_I is a zero column vector of dimension s , problem (2.2) is relaxed to a linear SDP:

$$(2.4) \quad \begin{aligned} & \text{minimize} && \sum_{(i,j) \in N_1} (\alpha_{ij}^+ + \alpha_{ij}^-) + \sum_{(i,k) \in N_2} (\alpha_{ik}^+ + \alpha_{ik}^-) \\ & \text{subject to} && \text{diag}(A_I^T Z A_I) = b_I, \\ & && \begin{pmatrix} e_{ij} \\ \mathbf{0} \end{pmatrix}^T Z \begin{pmatrix} e_{ij} \\ \mathbf{0} \end{pmatrix} - \alpha_{ij}^+ + \alpha_{ij}^- = (\widehat{d}_{ij})^2 \quad \forall (i,j) \in N_1, \\ & && \begin{pmatrix} e_i \\ -a_k \end{pmatrix}^T Z \begin{pmatrix} e_i \\ -a_k \end{pmatrix} - \alpha_{ik}^+ + \alpha_{ik}^- = (\widehat{d}_{ik})^2 \quad \forall (i,k) \in N_2, \\ & && \begin{pmatrix} e_{ij} \\ \mathbf{0} \end{pmatrix}^T Z \begin{pmatrix} e_{ij} \\ \mathbf{0} \end{pmatrix} \geq r_{ij}^2 \quad \forall (i,j) \in \overline{N}_1, \\ & && \begin{pmatrix} e_i \\ -a_k \end{pmatrix}^T Z \begin{pmatrix} e_i \\ -a_k \end{pmatrix} \geq r_{ik}^2 \quad \forall (i,k) \in \overline{N}_2, \\ & && Z \succeq 0, \quad \alpha_{ij}^+, \alpha_{ij}^-, \alpha_{ik}^+, \alpha_{ik}^- \geq 0, \\ & && i, j = 1, \dots, s, \quad k = s+1, \dots, n, \end{aligned}$$

where the constraint $\text{diag}(A_I^T Z A_I) = b_I$ ensures that the matrix variable Z 's lower right corner is a 2×2 identity matrix I , so that Z takes the form of Z_I in (2.3).

Initially, Biswas and Ye [2, 3] omit the \geq inequalities involving r_{ij} and r_{ik} and solve the resulting problem to obtain an initial solution Z_1 . (The inequality constraints increase the problem size dramatically, and Z_1 is likely to satisfy most of them.) They then adopt an ‘‘iterative active-constraint generation technique’’ in which inequalities violated by Z_k are added to the problem and the resulting SDP is solved to give Z_{k+1} ($k = 1, 2, \dots$). The process usually terminates before all constraints are included. Further study of this approach is presented in section 4.1.

2.4. SDP model analysis. Let $\bar{Z} = \begin{pmatrix} \bar{Y} & \bar{X}^T \\ \bar{X} & I \end{pmatrix}$ be a feasible solution of the relaxed SDP (2.4). Assuming the distance measurements are exact (no noise), Biswas and Ye [2] give conditions under which \bar{X} and \bar{Y} solve problem (2.2) exactly as follows:

- \bar{Z} is the unique optimal solution of (2.4), including all inequality constraints.
- In (2.4), there are $2n + n(n+1)/2$ exact pairwise distance measurements.

These conditions ensure that $\bar{Y} = \bar{X}^T \bar{X}$. In practice, distance measurements have noise and we only know that the SDP solution satisfies $\bar{Y} - \bar{X}^T \bar{X} \succeq 0$. This inequality can be used for error analysis of the location estimates provided by the relaxation. For example, $\text{trace}(\bar{Y} - \bar{X}^T \bar{X}) = \sum \tau_i$, where

$$(2.5) \quad \tau_i \equiv \bar{Y}_{ii} - \|\bar{x}_i\|^2 \geq 0,$$

is a measure of deviation of the SDP solution from the desired constraint $Y = X^T X$ (ignoring off-diagonal elements). The individual trace τ_i can be used to evaluate the location estimate \bar{x}_i for sensor i . In particular, we interpret a smaller τ_i to mean higher accuracy in the estimated location x_i . Further explanation is given in [2].

3. SpaseLoc: A scalable localization algorithm. When the number of nodes in (2.4) is large, applying a general SDP solver such as DSDP5.0 [1] or SeDuMi [23] would not scale well. In this section, we present a sequential subproblem approach named SpaseLoc to solve the full localization problem iteratively.

3.1. Adaptive subproblem approach. We call the overall sensor localization problem including all sensors and anchors the *full problem*. At each iteration, SpaseLoc selects from the *full problem* a subset of the unlocalized sensors and a subset of the anchors to form a localization *subproblem*. We call the selected sensors in the subproblem *subsensors*, and the selected anchors in the subproblem *subanchors*. These subsensors and subanchors, together with their known distance measurements and known anchors' locations, form a sub-SDP relaxation model to be solved using the same formulation as in (2.4).

In our adaptive approach, the subsensors and subanchors for each subproblem are chosen dynamically according to rule sets. (Rather than using predefined data, every new iteration's subproblem generation is based on the previous iteration's results.) The resulting SDP subproblems are of varying but limited size. Currently they are solved by Benson, Ye, and Zhang's SDP solver DSDP5.0 [1].

SpaseLoc is a *greedy algorithm* in the sense that each subproblem determines the final estimate of the associated sensor locations.

3.2. The SpaseLoc algorithm. The main steps of SpaseLoc are listed below, followed by explanations of the steps and definitions of new terms used therein.

- A0. Set *subproblem.size*.
- A1. Subproblem creation: Select subsensors and subanchors to be included in the subproblem.
- A2. Formulate SDP relaxation model (2.4) based on the chosen subsensors and subanchors, together with the known distances among them and the subanchors' known locations.
- A3. Call SDP solver to obtain a solution for the subsensors' locations.
- A4. Classify localized subsensors according to their τ_i value.
- A5. If all sensors in the network become localized or are determined to be outliers, go to step A6. Otherwise, return to step A1 for the next iteration.
- A6. Output all sensor locations and report outliers if any. Stop.




FIG. 3.1. *SpaseLoc* execution time as a function of *subproblem_size*: total nodes = 10000, anchors = 100, radius = 0.0226.

In step A0, *subproblem_size* specifies a limit on the number of unlocalized sensors to be included in each subproblem. It can range from 1 to an upper limit value that is potentially solvable by the SDP solver. In our experiments, the upper limit is 150. The most effective *subproblem_size* seems to change with the *full_problem* size, the model parameters such as *radius*, and the SDP solver used. We perform an approximate linesearch to find *subproblem_size* that corresponds to the minimum time because, empirically, the total execution time with all other parameters fixed is essentially a convex function of *subproblem_size*.

For example, when *full_problem* size is 10000 with 100 anchors, *radius* 0.0226, and no noise, *subproblem_size* 7 seems to give the best execution time with the DSDP5.0 solver (refer to Figure 3.1). The search time for *subproblem_size* is not included as part of the *SpaseLoc* execution time.

Step A1 involves choosing a subset of unlocalized sensors (no more than *subproblem_size*) and an associated subset of nodes with known locations. The latter can include a subset of the original anchors and/or a subset of sensors already localized by a previous subproblem (we define them as *acting anchors*). The rules for choosing subsensors and subanchors in this iteration are discussed in sections 3.4–3.5.

In step A4, the error in sensor i 's location is estimated by its individual trace τ_i as discussed in section 2.4. Subsensors whose τ_i value is within a given tolerance τ are labeled as localized and treated as acting anchors for the next iteration, whereas subsensors whose localization error is higher than the tolerance are also labeled as localized but are not used as acting anchors in later iterations. These new acting anchors are labeled with different acting levels as explained in section 3.4. The value of τ has an impact on the localization accuracy. Bigger values allow more localized sensors to be acting anchors, but with possibly greater transitive errors. Smaller values may increase the estimation accuracy for some of the sensors, but could lead to fewer connections to anchors for some unlocalized sensors. A rule of thumb is to use a small τ for networks with high anchor density to achieve potentially more accuracy, and a bigger τ for networks with low anchor density to avoid lacking connections to anchors. In order to avoid the side effect of a bigger τ eliminating too many potential acting anchors, at some later iteration we utilize all localized sensors as acting anchors

(including those whose τ_i value is bigger than the given tolerance τ). This change only starts when the remaining unlocalized sensors are connected to fewer than three anchors. It makes sure that we use acting anchors with higher accuracy first, but if no such acting anchors are available, we use localized points whose locations might be less accurate. In most cases, this is better than using no reference points at all.

In step A5, an unlocalized sensor is called an *outlier* when it does not have any distance information for the algorithm to decide its location. If a sensor has no connection to any anchor, it is classified as an outlier. In addition, if a connected cluster of sensors has no connection to any anchors, then all sensors in the cluster will be outliers.

The next sections explain the subproblem creation procedure used by step A1. Section 3.3 lists steps S1–S8 of the creation procedure itself. Section 3.4 presents rules RS1–RS4 for subsensor selection in step S5. Section 3.5 presents rules RA1–RA3 for subanchor selection in step S7. Section 3.6 illustrates the method for independent subanchor selection used in rules RA2–RA3. Sections 3.7–3.8 discuss the routines used in step S8 to localize sensors that have fewer than 3 connected anchors.

3.3. Subproblem creation procedure. As explained, *subproblem_size* is a predetermined parameter that represents the maximum number of unlocalized sensors that can be selected as subsensors in a subproblem. When there are more than *subproblem_size* unlocalized sensors, we have a choice to make among them.

The subproblem creation procedure makes sure that the choice of subsensors is based first on the number of connected anchors they have, and second on the type of connected anchors such as original anchors and different levels of acting anchors as defined in section 3.4, and that the choice of subanchors is based on a set of rules (section 3.5). The main steps are listed below, followed by explanations of the steps and definitions of new terms used.

- S1. Specify *MaxAnchorReq*.
- S2. Initialize *AnchorReq* = *MaxAnchorReq*.
- S3. Loop through unlocalized sensors, finding all that are connected to at least *AnchorReq* anchors. If *AnchorReq* ≥ 3 , determine if there are 3 *independent* subanchors; if not, go to the next sensor.¹ Enter each found sensor into a *candidate subsensor list*, and enter its connected anchors into a corresponding *candidate subanchor list*. Each sensor in the candidate subsensor list has its own candidate subanchor list (so there are as many candidate subanchor lists as the number of sensors in the candidate subsensor list). Let *sub_s_candidate* be the length of the candidate subsensor list.
- S4. If $0 < \textit{sub_s_candidate} \leq \textit{subproblem_size}$, the candidate subsensor list becomes the chosen subsensors list. Go to step S7.
- S5. If $\textit{sub_s_candidate} > \textit{subproblem_size}$, the choice of subsensors is further based on subsensor selection rules RS1–RS4 described in section 3.4. After exactly *subproblem_size* subsensors are selected from the candidate list according these rules, go to step S7.
- S6. Now *sub_s_candidate* = 0. Reduce *AnchorReq* by 1.
If *AnchorReq* > 0, go to step S3 for another round of subproblem creation.
Otherwise, *AnchorReq* = 0 and *sub_s_candidate* = 0 indicates that there are still unlocalized sensors left that are not connected to any localized node. We classify them as outliers and exit this procedure to continue at step A6 of

¹See section 3.6 for dependency definition and independent anchor selection.

section 3.2.

S7. Now that we have a subsensor list and the candidate subanchor lists, choose subanchors using selection rules RA1–RA3 presented in section 3.5.

S8. The subsensors and subanchors are selected and the subproblem creation routine finishes here.

If $AnchorReq \geq 3$, go to step A2 in section 3.2.

If $AnchorReq = 2$, apply the procedure in section 3.7 and go to step A5.

If $AnchorReq = 1$, apply the procedure in section 3.8 and go to step A5.

In step S1, $MaxAnchorReq$ determines the initial (maximum) value of $AnchorReq$. It is useful for scalability when connectivity is dense. A smaller $MaxAnchorReq$ would generally cause fewer subanchors to be included in the subproblem, thus reducing the number of distance constraints in each SDP subproblem and hence reducing execution time for each iteration. For instance, under ideal conditions (where there is no noise), even if a sensor has 10 distance measurements to 10 anchors, we don't need to include all 10 anchors because we can use 3 to localize that sensor accurately.

In the presence of noise, a bigger $MaxAnchorReq$ should reduce the average estimation error. For example, if there is a large distance measurement error from one particular anchor, since $MaxAnchorReq$ anchors are all taken into consideration for deciding the sensor's actual location, the large error would be averaged out. Another consideration for setting $MaxAnchorReq$ is the trade-off between estimation accuracy and execution speed. If we are in a static environment and would like to have localization as accurate as possible under noise conditions, we might choose a large $MaxAnchorReq$. However, in a real-time environment involving moving sensors, where speed might take priority, we would consider a smaller $MaxAnchorReq$.

In step S2, $AnchorReq$ is a dynamic parameter that may decrease in later steps.

In step S4, the subproblem may contain fewer than $subproblem_size$ subsensors, which is perfectly acceptable. The alternative is to reduce $AnchorReq$ by 1 and find more subsensor candidates that have fewer distance connections. However, this approach might reduce the accuracy of the algorithm, because we do want to localize the subsensors as accurately as possible as the iteration progresses, and the newly localized subsensors could be further used as acting anchors for the next iteration.

In step S6, $AnchorReq$ is iteratively reduced by 1 from $MaxAnchorReq$ to 0 eventually. This approach allows sensors with at least $AnchorReq$ connections to anchors to be localized before sensors with fewer connections to anchors.

As we know, under no-noise conditions, a sensor's location can be uniquely determined by connections to at least 3 independent anchors. If a sensor has connections to only 2 anchors, there are two possible locations; and if there is only 1 connection to an anchor, the sensor can be anywhere on a circle. In step S8, we use heuristic subroutines described in sections 3.7–3.8 to include the sensor's anchors' connected neighboring nodes in the subproblem in order to improve the estimation accuracy.

3.4. Subsensor selection. In step S5, when the number of sensors in the candidate subsensor list is bigger than $subproblem_size$, the choice of subsensors is further based on the types of anchors each sensor is connected to.

First, we introduce the concept of *sensor priority*. We assign a priority to each sensor in the candidate subsensor list. A sensor with a smaller priority value is selected to be localized before one with a bigger priority value. A sensor's priority is based on the types of anchors the sensor is directly connected to. Next, in order to define different types of anchors, we introduce the concept of *anchor acting levels*. All anchors including acting anchors are assigned certain acting levels. Original anchors

TABLE 3.1
An example: priority list when $MaxAnchorReq = 3$.

Priority value	Level 1 anchor	Level 3 anchor	Level 5 anchor	Level 7 anchor	Level 9 anchor	...	Resulting anchor level	
1	≥ 3	any						3
2	$= 2$	≥ 1	any					5
3	$= 1$	≥ 2	any					7
3	$= 2$	$= 0$	≥ 1	any				7
4	$= 1$	$= 1$	≥ 1	any				9
4	$= 2$	$= 0$	$= 0$	≥ 1	any			9
5	$= 1$	$= 1$	$= 0$	≥ 1	any			11
5	$= 2$	$= 0$	$= 0$	$= 0$	≥ 1	any		11
...	$= 0$	total ≥ 3						(11, $bigN$)
$bigN$				total = 2				$bigN$
$bigN+1$				total = 1				$bigN+1$

are always set to acting level 1. Every acting anchor is set to an acting level after it has been localized as a sensor. Essentially, acting anchors are set with acting levels depending on the levels of the anchors that localized them.

The priority rules for selecting subsensors from a candidate subsensor list are as follows:

- RS1. When $AnchorReq \geq 3$ and a sensor has at least 3 connected anchors that are independent, the sensor's priority depends on the 3 connected anchors that have the lowest acting levels among all its connected anchors. The sensor's priority value is defined as the summation of these 3 connected anchors' acting levels.
- RS2. If the sensor has 3 connected anchors that are dependent, it is ranked with the same priority as when the sensor is connected to only 2 anchors.
- RS3. Sensors with 2 anchor connections are ranked with equal priority, independent of the acting levels of the 2 connected anchors. (This can be easily expanded to be more granular according to the connected anchors' acting levels.) Sensors in this category are assigned lower priority than any sensors that have at least 3 independent anchor connections.
- RS4. Sensors with 1 anchor connection are ranked with equal priority, independent of the acting level of the connected anchor. (Again, this can be more granular according to the connected anchor's acting level.) Sensors in this category are assigned lower priority than any sensors that have at least 2 anchor connections.

Table 3.1 illustrates the priority list for an example where $MaxAnchorReq = 3$ and the sensor's priority is determined by the 3 anchors that have the lowest acting levels among all the sensor's connected anchors. We can certainly add more granularity by further classifying the acting levels of the sensor's fourth or fifth connected anchors (if any). Although more categorization of the priorities should increase localization accuracy under most noise conditions, more computational effort is required to handle more levels of priorities.

Each item in the table represents the number of anchors with different acting levels that are needed at each priority. The last column represents the resulting acting anchors' acting levels for subsequent iterations. For example, if a sensor has at least three independent connections to anchors, and if 3 of the anchors are original anchors (acting level 1), this sensor belongs to priority 1 as listed in row 1 of the table. Also, when this sensor is localized, it becomes acting anchor level 3 (the summation of the

anchor levels of the three anchors that localized it). Similarly, if a sensor has at least three independent connections to anchors, and if 2 of the anchors are original anchors (acting level 1) and at least 1 of the connected anchors is at acting level 3, then this sensor belongs to priority 2 as listed in row 2 of the table. Also, when this sensor is localized, it becomes acting anchor level 5. The sensors that connect to two anchors belong to the second to last priority, and sensors that connect to only one anchor belong to the last priority. We use a big enough number $bigN$ in the implementation to ensure that sensors connected to fewer than 3 anchors are given the lowest priority.

3.5. Subanchor selection. In step S7, for each unlocalized subsensor in the subsensor list, only *AnchorReq* of the connected anchors are allowed to be included in the subproblem. We use the following rules to select subanchors from a candidate subanchor list that contains more than *AnchorReq* anchors.

RA1. Original anchors are selected first, followed by acting anchors with lower acting level.

RA2. The subanchors chosen should be linearly independent.

RA3. Among independent anchors in the candidate subanchor list, we use distance scale-factors to encourage selection of the closest subanchors.

Rules RA2 and RA3 are implemented as in section 3.6. Rule RA3 is based on the assumption that under noise conditions, we trust the shorter distance measurements more than the longer ones.

3.6. Independent subanchors selection. Suppose sensor i is connected to K ($K > 3$) anchors at locations a_{ik} with corresponding distance measurements \hat{d}_{ik} ($k = 1, \dots, K$). Define the matrices

$$A = \begin{pmatrix} 1 & 1 & \dots & 1 \\ -a_{i1} & -a_{i2} & \dots & -a_{iK} \end{pmatrix}, \quad D_1 = \text{diag}(1/\sqrt{1 + \|a_{ik}\|^2}), \quad D_2 = \text{diag}(1/\hat{d}_{ik}).$$

We select an independent subset by a QR factorization with column interchanges [11]: $B = AD_1D_2$, $BP = QR$, where Q is orthogonal, R is upper-trapezoidal, and P is a permutation chosen to maximize the next diagonal of R at each stage of the factorization. (D_1 normalizes the columns of A , and D_2 biases them in favor of anchors that are closer to sensor i .) If the 3rd diagonal of R is larger than a predefined threshold (10^{-4} is used in our simulation), then the first 3 columns of AP are regarded as independent, and the associated anchors are chosen. Otherwise, all subsets of 3 among the K anchors are regarded as dependent. (In MATLAB, R and P are obtained by a command of the form `[Q,R,P] = qr(B)`.)

3.7. Geometric subroutine (two connected anchors). This section illustrates the heuristic techniques used in step S8 of section 3.3 to localize sensors connected to only two anchors.

When a sensor's connected anchors are also connected to other anchors, this subroutine may improve the accuracy of the sensor's localization, as illustrated by an example in Figure 3.2.

In this example, assume s_1 and s_2 are sensors with unknown locations, and $a_3(1, 3)$, $a_4(1, 2)$, $a_5(2, 2)$, $a_6(4, 1)$, $a_7(5, 1)$ are anchors with known locations in brackets. Assume that the sensors' radio range is $\sqrt{2}$, and we are also given two distance measurements $\hat{d}_{13} = 1$ and $\hat{d}_{14} = \sqrt{2}$ for sensor s_1 and one measurement $\hat{d}_{27} = 1$ for sensor s_2 .




FIG. 3.2. Sensors with connections to at most two anchors.




FIG. 3.3. (a) Sensor with two anchors’ circles intersecting. (b) Sensor with two anchors, a_2 ’s circle in a_3 ’s. (c) Sensor with two anchors’ circles disjoint.

Given two distances \hat{d}_{13} and \hat{d}_{14} to two anchors $a_3(1, 3)$ and $a_4(1, 2)$, we know that s_1 should be at either $(0, 3)$ or $(2, 3)$. If we only use $s_1, a_3(1, 3), a_4(1, 2)$ in an SDP subproblem, then SDP relaxation will give a solution near the middle of the two possible points, which would be very close to point $(1, 3)$. If there is any anchor (a_5) that is near s_1 ’s connected anchors (a_3 and a_4) with any of the two possible sensor points within their radio range (point $(2, 3)$ is within a_5 ’s range), that point $(2, 3)$ must not be the real location of s_1 , or else s_1 would be connected to this anchor (a_5) as well. Thus we can infer that s_1 must be at the other point $(0, 3)$.

Inspired by the above observation, when a sensor has at most 2 connected anchors, we include these anchors’ connected anchors in the subproblem (we call them the connected anchors’ neighboring anchors) together with the sensor and its directly connected anchors. By including the neighboring anchors, we might hope that the inequality constraints in the SDP relaxation model (2.4) would push the estimation towards the right point. However, because of the relaxation, enforcing inequalities in (2.4) is not equivalent to enforcing them in the distance model (2.2). The added inequality constraints only push the original solution near $(1, 3)$ a tiny bit towards s_1 ’s real location $(0, 3)$, and the solution essentially stays at around $(1, 3)$.

Given the ineffectiveness of the SDP relaxation approach under this condition, we propose instead a geometric approach as illustrated in Figure 3.3. Assume $s_1(x_x, x_y)$ has measurements \hat{d}_{12} to anchor $a_2(a_{2x}, a_{2y})$ and \hat{d}_{13} to anchor $a_3(a_{3x}, a_{3y})$. We also assume $\hat{d}_{12} \leq \hat{d}_{13}$ (we can always swap the two indexes otherwise). Let $a_l (l = 4, \dots, k)$ be a_2 and/or a_3 ’s neighboring anchors with radio range $r_{1l} (l = 4, \dots, k)$, and let d_{23} be the known (exact) Euclidean distance between a_2 and a_3 .

- If two circles centered at a_2 and a_3 with radii \hat{d}_{12} and \hat{d}_{13} intersect each other ($\hat{d}_{12} + \hat{d}_{13} \geq d_{23}$ and $\hat{d}_{13} - \hat{d}_{12} \leq d_{23}$) as in Figure 3.3(a):
 - Two possible locations of s_1 are given by solutions x^* and x^{**} of the

equations

$$\|x - a_2\|^2 = \widehat{d}_{12}^2, \quad \|x - a_3\|^2 = \widehat{d}_{13}^2.$$

- Sensor s_1 's location is selected from x^* and x^{**} , whichever is further away from any neighboring anchor. Thus, for $l = 4$ to k ,
 - if $\|x^* - a_l\|^2 < r_{1l}^2$, then $x = x^{**}$ and stop
 - else if $\|x^{**} - a_l\|^2 < r_{1l}^2$, then $x = x^*$ and stop.
 Otherwise, $x = (x^* + x^{**})/2$ and stop.
- Under noise conditions, the a_2 circle may be inside the a_3 circle ($\widehat{d}_{12} + \widehat{d}_{13} \geq d_{23}$ and $\widehat{d}_{13} - \widehat{d}_{12} > d_{23}$) as in Figure 3.3(b).
 - The solutions x^* and x^{**} of the following equations give two possible points for s_1 on the a_2 circle:

$$\begin{aligned} (x_x - a_{2x})^2 + (x_y - a_{2y})^2 &= \widehat{d}_{12}^2, \\ (a_{2x} - a_{3x})(x_y - a_{2y}) &= (a_{2y} - a_{3y})(x_x - a_{2x}), \end{aligned}$$

where x is on the line through a_2 and a_3 represented by the second equation.

- If $\|x^* - a_3\| < \|x^{**} - a_3\|$, then $x = x^{**}$; otherwise $x = x^*$. This guarantees that the point further from a_3 is chosen. Note that we base the sensor's estimation on the closest anchor (a_2 here since $\widehat{d}_{13} \geq \widehat{d}_{12}$), assuming that a shorter measurement is generally more accurate than longer ones, given similar anchor properties.

The same approach applies when the a_3 circle is inside the a_2 circle ($\widehat{d}_{12} - \widehat{d}_{13} > d_{23}$).

- Under noise conditions, the a_2 and a_3 circles may again have no intersection ($\widehat{d}_{12} + \widehat{d}_{13} < d_{23}$) as in Figure 3.3(c).
 - The solutions x^* and x^{**} of the following equations give two possible points for s_1 on the circle for the anchor with smaller *radius*. Let us assume $\widehat{d}_{12} \leq \widehat{d}_{13}$:

$$\begin{aligned} (x_x - a_{2x})^2 + (x_y - a_{2y})^2 &= \widehat{d}_{12}^2, \\ (a_{2x} - a_{3x})(x_y - a_{2y}) &= (a_{2y} - a_{3y})(x_x - a_{2x}), \end{aligned}$$

where x is on the line through a_2 and a_3 represented by the second equation.

- If $\|x^* - a_3\| > \|x^{**} - a_3\|$, then $x = x^{**}$; otherwise $x = x^*$. This guarantees that the point closer to a_3 (in between a_2 and a_3) is chosen.

3.8. Geometric subroutine (one connected anchor). Similar inefficiency occurs in the SDP solution when a sensor connects to only *one* anchor. The SDP solver under this condition gives a solution for the sensor to be in the same location as the sensor's connected anchor. In reality, the sensor could be anywhere on the circle. The SDP gives an average point, at the center of the circle, and that is where the connected anchor is. Even if the anchor's neighboring anchor is included in the SDP subproblem, the inequality constraints are not active most of the time because the SDP solution may not provide optimal solutions all the time.

We propose a heuristic for estimating a sensor's location with only one connecting anchor. The idea is to use one neighboring anchor's radio range information to



FIG. 3.4. (a) Sensor with one anchor connection a and one neighboring anchor b . (b) Sensor with one anchor connection a and two neighboring anchors b, c .

eliminate the portion of the circle that the sensor would not be on, and then calculate the middle of the other portion of the circle to be the sensor's location. For the example in Figure 3.2, because we know the distance between s_2 and a_7 is 1, we know that s_2 could be anywhere on the circle surrounding a_7 with *radius* 1. Knowing a_7 's neighboring anchor node a_6 is not connected to s_2 , we know that s_2 would not be in the area surrounding a_6 with *radius* $\sqrt{2}$. Thus, s_2 could be anywhere around the half circle including points (5, 2), (6, 1), (5, 0). The heuristic chooses the middle point between the two circles' intersection points (5, 2) and (5, 0), which happens to be (6, 1) in this example. The heuristic gives better accuracy for the sensor's location than the SDP solution under most conditions. The procedure follows.

- Assume s has one distance measurement \hat{d} to anchor a , and b is the closest connected neighboring anchor to a with radio range r (refer to Figure 3.4(a)). We assume $a = (a_x, a_y)$, $b = (b_x, b_y)$, $x = (x_x, x_y)$.
- The solutions x^* and x^{**} of the following equations give two possible points s on the circle:

$$\begin{aligned} (x_x - a_x)^2 + (x_y - a_y)^2 &= \hat{d}^2, \\ (a_x - b_x)(x_y - a_y) &= (a_y - b_y)(x_x - a_x), \end{aligned}$$

where x is on the line through a and b represented by the second equation.

- If $\|x^* - b\| < r$, then $x = x^{**}$; otherwise $x = x^*$. This guarantees that the point further from b is chosen.

The above heuristic provides a simple way of estimating a sensor's location when the sensor connects to only one anchor. A more complicated approach can be adopted when the connected anchor has more than one neighboring anchor, which can increase the accuracy of the sensor's location. We call it an arc elimination heuristic. The idea is to loop through each of the neighboring anchors and find the portion of the circle that the sensor won't be on, and eliminate that arc as a possible location of the sensor. Eventually, when one or more plausible arcs remain, we choose the middle of the largest arc to be the sensor's location. For example, assume we add one more neighboring anchor c to sensor s 's anchor a from the previous example in Figure 3.4(a). The new scenario is shown in Figure 3.4(b). First, we find the intersections (points 1 and 2) of two circles: one at a with *radius* \hat{d} , the other at b with *radius* r . We know that the 1-2 portion of the arc closer to point b won't be the location of s . Second, we find the intersections (points 3 and 4) of two circles: one at a with *radius* \hat{d} , the other at c with *radius* r . We know that the arc 3-4 closer to point c won't be the location of s . Thus we deduce that s must be somewhere on the arc 1-4 further away from b or c . The estimation of s is given in the middle of the arc 1-4. As we see, this method should provide more accuracy than the one-neighboring-anchor approach.

3.9. Subproblem optimality. For the case of one sensor connected to three independent anchors, Biswas and Ye [2] prove when there is no noise that the SDP relaxation (2.4) gives an optimal solution to (2.2). The proof depends on the fact that there are three independent equations and only three variables.

In SpaseLoc, the subproblems are constructed from sensors that have three independent anchors (or acting anchors) where possible. If each of these subsensors (say total s) were included in separate subproblems together with their connected 3 independent anchors, the proof in [2] shows that they would be localized exactly by the SDP approach. If these s subsensors and their connected anchors are treated together in a single subproblem, the larger SDP relaxation contains sets of the same three equations that would arise in the separate SDP relaxations. The equations form a block-diagonal system in the larger SDP. There are $3s$ independent equations and the same number of variables containing only x_i and y_{ii} , $i = 1, \dots, s$. The \widehat{d}_{ik} equation in (2.4) reduces under no noise conditions to $y_{ii} = x_i^T x_i$ for all relevant pairs (i, k) . The constraint $Y - X^T X \succeq 0$ then guarantees $y_{ij} = x_i^T x_j$ for all $j = 1, \dots, s$. Hence, the SDP solution for the SpaseLoc subproblem is also rank 2 and gives an exact locations for all subsensors.

4. Computational results. This section explains the simulation method and the setup for experimenting with the SpaseLoc algorithm, then presents results for various parameter settings.

For the simulation, a total number of nodes n (including s sensors and m anchors) is specified in the range 4 to 10000. The locations of these nodes are assigned with a uniform random distribution on a square region of size $r \times r$ where $r = 1$, or put on the grid-points of a regular topology such as a square or an equilateral triangle on the same region. The m anchors are randomly chosen from the given n nodes. We assume all sensors have the same radio range (*radius*) for any given test case. Various radio ranges were tested in the simulation.

Euclidean distances $d_{ij} = \|x_i - x_j\|$ are calculated among all sensor pairs (i, j) for $i < j$. We then use \widehat{d}_{ij} to simulate measured distances, where \widehat{d}_{ij} is d_{ij} times a random error simulated by *noise_factor* $\in [0, 1]$. For a given *radius* $\subseteq [0, 1]$ it is defined as follows:

- If $d_{ij} \leq \text{radius}$, then $\widehat{d}_{ij} = d_{ij}(1 + \text{rn} * \text{noise_factor})$, where *rn* is normally distributed with mean zero and variance one. (Any numbers generated outside $(-1, 1)$ are regenerated.)

In practical networks, depending on the technologies that are being used to obtain the distance measurements, there may be many factors that contribute to the noise level. For example, one way to obtain the distance measurement is to use the received radio signal strength between two sensors. The signal strength could be affected by media or obstacles in between the two sensors. In this study, *noise_factor* is a normally distributed random variable with mean zero and variance one. This model could be replaced by any other noise model in practice.

- If $d_{ij} > \text{radius}$, then $\widehat{d}_{ij} = 0$, and the bound $r_{ij} = 1.001 * \text{radius}$ is used in the SDP model.

In the simulation, we define the average estimation error to be $\frac{1}{s} \sum_{i=1}^s \|\bar{x}_i - x_i\|$, where \bar{x}_i is from the SDP solution and x_i is the i th node's true location. In a practical setting, we wouldn't know the node's true location x_i . Instead, we would use the node's trace τ_i (2.5) to gauge the estimation error.

To convey the distribution of estimation errors and trace, we also give the 95%

quartile.

Factors such as noise level, radio range, and anchor densities can directly impact localization accuracy. The sensors' estimated locations are derived directly from the given distance measurements. If the noise level in these measurements is high, the estimation accuracy cannot be high. We also need sufficiently large radio range to achieve accurate localization, because too small a range could cause many sensors to have low connectivity or even be unreachable. Finally, more anchors in the network should help with the estimation accuracy because there are more reference points.

In the following subsections, we present simulation results (most results averaged over 10 runs) to show the accuracy and scalability of the SpaseLoc algorithm. We observe the impact of various radio ranges, anchor densities, and noise levels on the accuracy and performance of the algorithm. Computations were performed on a laptop computer with 2.4 GHz CPU and 1GB system memory, using MATLAB 6.5 [16] for SpaseLoc and a Mex interface to DSDP5.0 (Benson, Ye, and Zhang [1]) for the SDP solutions.

4.1. Effect of inequality constraints in SDP relaxation model. As we discussed in section 3.7, because of the $Y \succeq X^T X$ constraint relaxation, enforcing the r_{ij}^2 and r_{ik}^2 inequality constraints in (2.4) is not equivalent to enforcing them in the distance model (2.2). In order to observe the effectiveness of including these inequality constraints, we conduct simulations with the following three strategies, according to the number of times we check for violated inequality constraints and then include them to obtain a new solution.

- I0. This corresponds to solving the SDP problem with equality constraints only. (No inequality constraints are ever added.) The final solution is optimal for problem (2.4) without the inequality constraints involving r_{ij}^2 and r_{ik}^2 .
- I1. This corresponds to solving the SDP problem with all equalities (and no inequalities) first, and then adding violated inequality constraints and resolving it at most once.
- I2. This corresponds to solving the SDP problem with all equalities (and no inequalities) first, and then adding violated inequality constraints and resolving one or more times until all inequalities are satisfied. The final solution is an optimal solution to problem (2.4).

Our experimental results show that the added inequality constraints do not always provide better localization accuracy, but can greatly increase the execution time. In this section, we illustrate the inequality constraints' impact through two simulation examples: one with no noise but low connectivity; the other with full connectivity but with noise.

In our first example, we run simulation results on a network of 100 randomly uniform-distributed sensors with *radius* 0.2275 and 10 randomly selected anchors. One of the sensors happens to be connected to only two other nodes. The sensors are localized with the full SDP and with SpaseLoc, using each of the I0, I1, I2 strategies in turn. In addition, we examine each case with or without our geometric routines for SpaseLoc. The results are shown in Figure 4.1 and Table 4.1.

Figure 4.1 shows there is a sensor connected to only 2 anchors. For full SDP shown in (a), no violated inequalities are ever found, so full SDP with I0, I1, or I2 has only one SDP call and always generates the same results. For SpaseLoc in (b) with I0 and no geometric routine, SDP is called 47 times (with no subsequent check for violated constraints). It produces the similar estimation accuracy as the full SDP approach but with much improved performance. In (c), SpaseLoc with I1 or

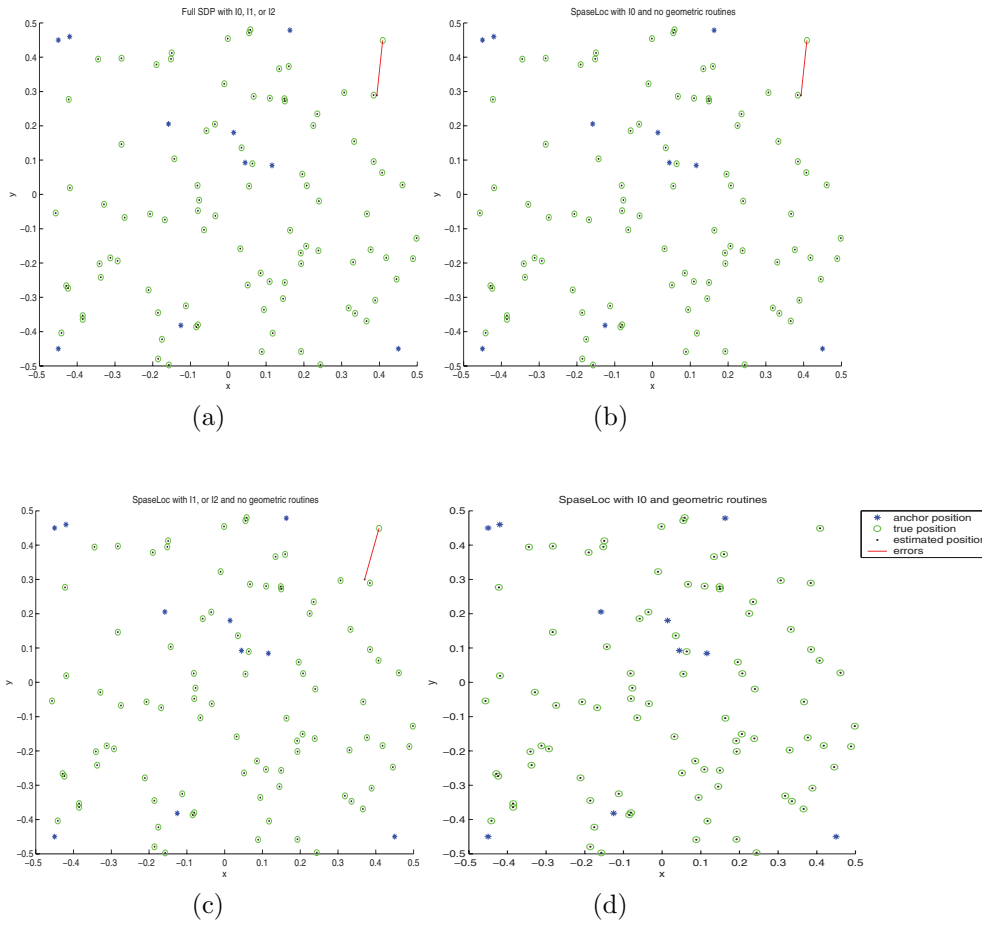


FIG. 4.1. Inequality impact on accuracy: 100 nodes, 10 anchors, no noise, radius 0.2275.

TABLE 4.1

Inequality impact on accuracy and speed: 100 nodes, 10 anchors, no noise, radius 0.2275.

Methods	Error	95% Error	Time	SDP's
Full SDP with I0 or I1 or I2	1.7877e-3	1.7483e-10	11.97	1
SpaseLoc with I0 and no geometric routines	1.7890e-3	1.1684e-7	0.38	47
SpaseLoc with I1 or I2 and no geometric routines	1.7134e-3	1.1684e-7	0.42	48
SpaseLoc with I0 and geometric routines	1.4679e-7	1.1523e-7	0.35	46

I2 produces the same results, which means violated inequalities are found only once. Comparing (b) and (c), we see that including violated inequalities does improve the estimation accuracy a little. Best of all in (d), SpaseLoc with I0 and our geometric routines localizes all sensors with virtually no error.

Table 4.1 shows that adding violated inequalities increases execution time slightly for SpaseLoc.

In our second example, in order to observe the effectiveness of the inequality constraints under noise conditions, we run simulations for a network of 100 nodes whose true locations are at the vertices of an equilateral triangle grid. Ten anchors are placed at the middle grid-point of each row, and the *radius* is 0.25. A *noise_factor*




FIG. 4.2. *Inequality impact on accuracy: 100 nodes, 10 anchors, noise_factor 0.1, radius 0.25.*

TABLE 4.2

Inequality impact on accuracy and speed: 100 nodes, 10 anchors, noise_factor 0.1, radius 0.25.

Methods	Error	Time	SDP's
Full SDP with I0	0.1268	13.87	1
Full SDP with I1	0.1292	34.20	2
Full SDP with I2	0.1403	134.50	4
SpaseLoc with I0	0.0231	0.45	54
SpaseLoc with I1 or I2	0.0203	0.51	56

of 0.1 is applied to the distance measurements. The sensors are localized with either full SDP or SpaseLoc using I0, I1, I2 in turn without geometric routines. (Although we do not activate the geometric routines in this experiment, they are not a factor here because the localization error is not caused by low connectivity but by the noisy measurements.) The results are shown in Figure 4.2 and Table 4.2. Figure 4.2(b) and (d) correspond to strategy I1 or I2 for full SDP and SpaseLoc.

As we can see, adding violated inequalities for full SDP not only increases the execution times dramatically, but also increases the localization error. For SpaseLoc, adding violated inequalities improves the estimation accuracy slightly. Note that I1

and I2 produce the same results for SpaseLoc.

In summary, the first experiment shows that when the errors are caused by *low connectivity*, SpaseLoc with geometric routines and no inequality constraints (I0) outperforms SpaseLoc with inequalities (I1 or I2) and all of the full SDP options. Given this observation, from now on we only use SpaseLoc with geometric routines, which means the geometric routines are used instead of SDP to localize sensors connected to less than 3 anchors.

The second experiment indicates that under noise conditions, although adding violated inequalities does not seem to improve the estimation accuracy for full SDP, it does improve accuracy for SpaseLoc.

In the subsequent sections, we continue to examine the inequality constraints' effects on accuracy and speed.

4.2. Accuracy and speed comparison: Full SDP versus SpaseLoc. For very small networks, the SDP approach is both accurate and efficient. (This is vital to SpaseLoc, as many small subproblems must be solved using SDP.) However, the performance of the full SDP approach deteriorates rapidly with network size.

Table 4.3 shows the localization results using full SDP (a) and using SpaseLoc (b) for a range of examples with various numbers of nodes whose true locations in the network are at the vertices of an equilateral triangle grid. Anchors are placed at the middle grid-point of each row. A *noise_factor* of 0.1 is applied to the distance measurements.

Let us first look at the impact of I0, I1, and I2 on estimation accuracy. Table 4.3 (a) shows that for full SDP, 8 errors with I1 are bigger than with I2, and 5 errors with I2 are bigger than with I1. Comparing I0 with I2, we see that for each strategy, 9 errors in I0 are bigger than the errors for the other strategy. It appears that full SDP with added inequalities does not always improve the estimation accuracy. For SpaseLoc, I1 and I2 generate almost equivalent estimation accuracy; I0 has 8 errors that are bigger than with I1, while I1 has 5 errors bigger than with I0. Therefore, the added inequalities provide only marginal accuracy improvement for SpaseLoc.

Now let us compare full SDP with SpaseLoc. Figures 4.3–4.4 plot results for full SDP with I0 and SpaseLoc with I0 for two of these examples: 9 and 49 nodes, including 3 and 7 anchors placed at the grid-point in the middle of each row. As we can see from these two figures and Table 4.3, for localizing 4 and 9 nodes, full SDP and SpaseLoc show comparable performance. Beyond that size, the contrast grows rapidly. For localizing 49 nodes, SpaseLoc is 10 times faster than the full SDP method, with more than four times the accuracy. For 400 nodes, SpaseLoc with strategies I0, I1, and I2 is, respectively, 800, 2500, and 8500 times faster than full SDP with the same strategies, while achieving 10 times greater accuracy. Thus, the full SDP model becomes less effective as problem size increases. In fact, for problem sizes above 49 nodes, the average estimation error using full SDP becomes so large that the computed solution is of little value.

It may seem nonintuitive that SpaseLoc's greedy approach could produce smaller errors than the full SDP method. However, all of the SDP problems and subproblems of the form (2.4) are *relaxations* of Euclidean models of the form (2.2). As we discussed in section 3.9, SpaseLoc always tries to create a subproblem whose subsensors have three independent anchor connections, so that the SDP solution is exact. The same conclusion cannot be drawn under noise conditions, but experimentally the relaxations under noise conditions appear to be *tighter* in SpaseLoc's subproblems than in the single large SDP.

TABLE 4.3
Accuracy and speed comparison between full SDP and SpaseLoc.

(a) Full SDP

Number of nodes	Radio range	Error			Time (sec)			SDP calls		
		I0	I1	I2	I0	I1	I2	I0	I1	I2
4	2.24	0.0317	0.0317	0.0317	0.01	0.01	0.01	1	1	1
9	1.12	0.1267	0.1203	0.1203	0.02	0.05	0.05	1	2	2
16	0.75	0.0837	0.0703	0.0680	0.10	0.21	0.35	1	2	3
25	0.56	0.0938	0.1170	0.1170	0.37	0.80	1.26	1	2	3
36	0.45	0.0719	0.0618	0.0561	0.81	1.88	3.02	1	2	3
49	0.40	0.1190	0.1190	0.1190	2.10	5.33	5.33	1	2	2
64	0.40	0.1218	0.0919	0.0954	3.43	9.21	21.60	1	2	4
81	0.40	0.1380	0.0894	0.0885	7.26	19.66	59.05	1	2	5
100	0.25	0.1268	0.1292	0.1403	13.87	34.20	140.26	1	2	4
121	0.40	0.1157	0.1088	0.1091	23.24	81.62	182.74	1	2	3
144	0.21	0.1480	0.1899	0.1891	37.76	168.43	584.23	1	2	4
169	0.40	0.1283	0.1141	0.1217	71.87	278.72	692.12	1	2	4
196	0.18	0.1404	0.1275	0.1286	151.52	461.97	1081.35	1	2	4
225	0.40	0.1568	0.1589	0.1571	232.31	752.75	2408.67	1	2	5
256	0.15	0.1429	0.1375	0.1370	356.86	1089.52	3260.33	1	2	5
324	0.14	0.1685	0.1685	0.1685	962.66	2620.20	2620.20	1	2	2
361	0.13	0.1734	0.1842	0.1833	1391.04	5051.05	15281.26	1	2	4
400	0.12	0.1819	0.1970	0.1968	1662.22	5950.34	20321.60	1	2	4

(b) SpaseLoc

Number of nodes	Radio range	Error			Time (sec)			SDP calls		
		I0	I1	I2	I0	I1	I2	I0	I1	I2
4	2.24	0.0317	0.0317	0.0317	0.02	0.02	0.02	1	1	1
9	1.12	0.0513	0.0513	0.0513	0.04	0.04	0.04	6	6	6
16	0.75	0.0615	0.0559	0.0559	0.06	0.19	0.09	8	9	9
25	0.56	0.0597	0.0608	0.0608	0.13	0.13	0.13	12	13	13
36	0.45	0.0364	0.0294	0.0294	0.17	0.20	0.20	20	23	23
49	0.40	0.0252	0.0252	0.0252	0.21	0.21	0.21	26	26	26
64	0.40	0.0272	0.0273	0.0273	0.30	0.34	0.34	38	42	42
81	0.40	0.0286	0.0295	0.0295	0.37	0.41	0.41	49	53	53
100	0.25	0.0232	0.0203	0.0203	0.46	0.49	0.49	54	56	56
121	0.40	0.0238	0.0227	0.0227	0.57	0.61	0.61	74	77	77
144	0.21	0.0230	0.0237	0.0237	0.69	0.70	0.70	84	89	89
169	0.40	0.0200	0.0190	0.0190	0.80	0.84	0.84	100	106	106
196	0.18	0.0177	0.0177	0.0177	0.98	1.08	1.08	84	90	90
225	0.40	0.0226	0.0207	0.0208	1.07	1.41	1.47	94	109	110
256	0.15	0.0208	0.0235	0.0249	1.21	1.44	1.50	118	131	132
324	0.14	0.0179	0.0178	0.0178	1.64	1.70	1.70	157	158	158
361	0.13	0.0218	0.0217	0.0217	1.89	2.01	2.01	177	181	181
400	0.12	0.0176	0.0175	0.0175	2.02	2.37	2.37	184	201	201

In the following sections, we run more simulations only with SpaseLoc.

4.3. Scalability. Table 4.4 shows simulation results for 49 to 10000 randomly uniform-distributed sensors being localized using SpaseLoc with strategies I0, I1, and I2. The node numbers 49, 100, 225, ... are squares k^2 , and the *radius* is the minimum value that permits localization on a regular $k \times k$ grid. The number of anchors changes with the number of sensors and is chosen to be k . Noise is not included in this simulation. When the items under I1, I2 are empty, it means that they are equal to the values under I0 in the same row.




FIG. 4.3. 9 nodes on equilateral-triangle grids, 3 anchors, 0.1 noise, radius 1.12.





FIG. 4.4. 49 nodes on equilateral-triangle grids, 7 anchors, 0.1 noise, radius 0.40.

TABLE 4.4
SpaseLoc scalability. Strategies I1 and I2 generate same results.

Nodes	An_ chors	Ra_ dius	Sub_ size	Error		95% Error	Time		SDP's	
				I0	I1,I2	I0,I1,I2	I0	I1,I2	I0	I1,I2
49	7	0.3412	3	4.5840e-8		3.4449e-8	0.18		18	
100	10	0.2275	3	1.4679e-7		1.1523e-7	0.35		46	
225	15	0.1462	3	4.4940e-7		3.1248e-7	0.82		112	
529	23	0.0931	3	2.1662e-6		8.9873e-7	2.02		278	
1089	33	0.0620	3	1.1969e-4		7.1510e-5	4.48		587	
2025	45	0.0451	4	1.4917e-4	1.4115e-4	9.6639e-5	8.85	9.28	1006	1007
3969	63	0.0334	4	1.2399e-4		7.2414e-5	18.79		1867	
5041	71	0.0319	6	1.5172e-4		1.1918e-4	27.19		2210	
6084	78	0.0290	6	1.7126e-4		1.1475e-4	33.66		2742	
7056	84	0.0269	7	5.2369e-5		4.0388e-5	40.59		3117	
8100	90	0.0251	7	2.7376e-4	2.7353e-4	1.7071e-4	47.87	49.71	3564	3566
9025	95	0.0238	7	2.1141e-4	2.1977e-4	1.6039e-4	54.41	56.03	3957	3958
10000	100	0.0226	7	2.0269e-4		1.5836e-4	59.33		4452	

We find that strategies I1 and I2 produce the same results, and I0 gives essentially the same. This is because the inaccuracy of the estimation is caused purely by low connectivity, not by noisy distance measurements. Empirically we see that the program scales well: almost linearly in the number of nodes in the network. Indeed, the computational complexity of the SpaseLoc algorithm is of order n , the number of

FIG. 4.5. *SDP computational complexity.*

sensors in the network, even though the full SDP approach has much greater complexity, as we now show.

We know that in the full SDP model (2.4), the number of constraints is $O(n^2)$, and in each iteration of its interior-point algorithm the SDP solver needs to solve a *sparse* linear system of equations whose dimension is the number of constraints. Figure 4.5 plots the CPU time for strategy I0 from Table 4.3(a) as well as three curves of the form $time = a_p n^p$ for $p = 2, 3, 4$, where a_p is determined by a least-squares fit. It appears that the SDP complexity with strategy I0 lies somewhere between $O(n^3)$ and $O(n^4)$.

In SpaseLoc, we partition the full problem into p subproblems of size q or less, where $p \times q = n$. We generally set q to be much smaller than n , ranging from 2 to around 10 in most of our simulations. If t represents the execution time taken by the full SDP method for a 10-node network, in the worst case the computation time for SpaseLoc is $t \times O(p)$. Thus, SpaseLoc is really linear in p in theory. Since we can assume q to be a parameter ranging from 2 to 10, with worst case 2, we know that $O(p) = O(n/q) \leq O(n/2) = O(n)$. Now we can see that SpaseLoc's computation time is $O(n)$.

In the remaining subsections we choose the middle network size from Table 4.4 (nodes = 3969) to observe the effect of varying radio range, number of anchors, and noise.

4.4. Radio range impact. With a fixed total number of randomly uniform-distributed nodes (3969, of which 63 are anchors), Table 4.5 shows the direct impact of *radius* in the range 0.0304 to 0.0334 on accuracy and performance.

Strategies I1 and I2 produce essentially the same results, and with slightly better accuracy than I0 for 8 of the 16 *radius* values, while I0 produces slightly better accuracy than I1 or I2 in 4 cases. However, I1 and I2 take more time than I0 because they need more SDP calls.

TABLE 4.5

Radio range impact: nodes = 3969, anchors = 63, no noise, sub_size = 5.

radius	Error			95% Error			Time			SDP's		
	I0	I1	I2	I0	I1	I2	I0	I1	I2	I0	I1	I2
0.0304	2.444e-3	2.359e-3	2.359e-3	4.035e-4	5.757e-4	5.757e-4	18.03	18.60	18.60	1743	1688	1689
0.0306	1.122e-3	1.123e-3	1.123e-3	5.638e-4	5.644e-4	5.644e-4	18.13	18.70	18.70	1747	1749	1749
0.0308	2.460e-3	1.412e-3	1.412e-3	7.952e-4	6.039e-4	6.039e-4	18.64	19.39	19.51	1879	1895	1896
0.0310	1.087e-3	1.083e-3	1.083e-3	4.424e-4	4.397e-4	4.397e-4	18.39	19.05	19.05	1809	1814	1814
0.0312	2.480e-3	2.481e-3	2.481e-3	3.142e-4	3.146e-4	3.146e-4	18.30	18.93	18.93	1715	1717	1717
0.0314	5.464e-4	5.337e-4	5.337e-4	2.612e-4	2.612e-4	2.612e-4	18.90	19.53	19.53	1897	1900	1900
0.0316	4.828e-4	4.827e-4	4.827e-4	2.645e-4	2.645e-4	2.645e-4	18.95	19.54	19.54	1916	1917	1917
0.0318	3.018e-4	3.013e-4	3.013e-4	1.955e-4	1.955e-4	1.955e-4	19.06	19.65	19.65	1911	1913	1913
0.0320	4.214e-4	4.214e-4	4.214e-4	1.781e-4	1.781e-4	1.781e-4	18.91	19.49	19.49	1847	1848	1848
0.0322	2.842e-4	2.842e-4	2.842e-4	1.702e-4	1.702e-4	1.702e-4	18.89	19.45	19.45	1894	1895	1895
0.0324	5.213e-4	5.495e-4	5.495e-4	2.968e-4	3.020e-4	3.020e-4	18.91	19.58	19.71	1859	1865	1866
0.0326	4.091e-4	4.033e-4	4.033e-4	2.323e-4	2.315e-4	2.315e-4	18.96	19.51	19.51	1890	1893	1893
0.0328	2.299e-4	2.289e-4	2.289e-4	1.363e-4	1.363e-4	1.363e-4	18.87	19.46	19.46	1921	1922	1922
0.0330	2.057e-4	2.160e-4	2.161e-4	9.435e-5	9.450e-5	9.450e-5	18.83	19.52	19.64	1873	1875	1876
0.0332	6.192e-4	6.439e-4	6.439e-4	3.557e-4	3.557e-4	3.557e-4	19.37	20.04	20.04	1849	1853	1853
0.0334	1.240e-4	1.240e-4	1.240e-4	7.241e-5	7.241e-5	7.241e-5	18.79	18.79	18.79	1867	1867	1867

TABLE 4.6

Number of anchors impact: nodes = 3969, radius = 0.0334, no noise, sub_size = 5.

Anchors	Error		95% Error		Time		SDP's	
	I0	I1,I2	I0	I1,I2	I0	I1,I2	I0	I1,I2
40	1.052e-3	1.052e-3	8.408e-4	8.409e-4	19.30	19.88	1906	1908
50	1.109e-3	1.128e-3	7.748e-4	7.745e-4	19.38	20.15	1861	1865
100	8.782e-4	7.280e-4	5.337e-4	5.115e-4	19.16	19.96	1870	1872
150	2.716e-4	2.717e-4	1.025e-4	1.025e-4	18.86	19.60	1806	1808
200	4.889e-5	4.872e-5	1.473e-5	1.473e-5	18.77	19.52	1795	1796
250	1.716e-5	1.699e-5	7.760e-6	7.760e-6	18.55	19.32	1748	1749
300	1.538e-5	1.521e-5	4.408e-6	4.408e-6	18.20	18.99	1750	1751
350	7.533e-6	7.365e-6	2.858e-6	2.858e-6	18.13	18.93	1684	1685
400	6.383e-6	6.215e-6	1.841e-6	1.841e-6	18.16	18.96	1560	1591

As we see, increasing *radius* leads to increased accuracy and only slightly more computational time. The simulation could assist sensor network designers in selecting a radio range to achieve a desired estimation accuracy with little concern about algorithm speed.

4.5. Number of anchors impact. With constant *radius* (0.0334) and the same randomly distributed nodes (3969), Table 4.6 shows the impact of the number of anchors, ranging from 1% to 10% of the total number of points. (Noise is not included.)

Strategies I1 and I2 produce identical results. Comparing I0 with I1 or I2, we see that added inequalities slightly improve the average error consistently, although the 95% error remains essentially the same. Increasing the number of anchors in the network improves the estimation accuracy in general, with no obvious impact on algorithm speed. However, we don't see accuracy improvement when the number of anchors reaches more than 10% of the total points. This analysis is beneficial for designers to avoid the cost of deploying unnecessary anchors.

4.6. Noise impact. With constant *radius* (0.0334) and the same randomly distributed nodes (3969), Table 4.7 shows the impact of noise conditions on accuracy and performance.

We see that strategies I1 and I2 do not provide consistent improvement over I0 for both average and 95% error, yet they always increase execution time. Also, more noise in the network has a direct impact on estimation accuracy. Simulations of this kind may help designers determine the measurement noise level that will give an acceptable estimation error.

TABLE 4.7
Noise_factor impact: nodes = 3969, anchors = 400, radius = 0.0334, sub_size = 5.

Noise factor	Error			95% Error			Time			SDP's		
	I0	I1	I2	I0	I1	I2	I0	I1	I2	I0	I1	I2
0.01	9.60e-4	9.59e-4	9.59e-4	1.90e-6	3.16e-4	3.16e-4	20.38	21.18	21.20	1622	1623	1623
0.05	3.15e-3	8.33e-3	8.33e-3	3.16e-4	3.57e-3	3.57e-3	21.56	21.75	21.86	1479	1253	1256
0.10	6.87e-3	7.36e-3	1.03e-2	1.91e-3	5.17e-3	6.95e-3	21.46	24.73	22.94	1447	1592	1230
0.20	1.55e-2	1.57e-2	1.65e-2	4.95e-3	1.16e-2	1.25e-2	21.55	25.67	25.83	1208	1433	1390
0.30	1.51e-2	1.48e-2	1.46e-2	1.17e-2	1.27e-2	1.24e-2	21.09	29.76	31.78	1411	1829	1844
0.40	1.98e-2	1.79e-2	1.79e-2	1.32e-2	1.57e-2	1.57e-2	21.30	32.05	35.15	1523	1985	2073
0.50	3.05e-2	2.35e-2	2.28e-2	2.86e-2	2.16e-2	2.08e-2	22.07	35.27	39.26	1608	2157	2252

5. Summary and extensions. We have shown that SpaseLoc achieves the aims of accuracy, speed, and scalability with a single processor on very large networks. It takes full advantage of the recent SDP approach of Biswas and Ye [2]. The latter has computational complexity $O(n^p)$, where n is the network size and p is between 3 and 4, but we use it on multiple tiny subproblems to obtain an algorithm with essentially linear complexity. On a 2.4GHz laptop with 1GB memory, SpaseLoc maintains efficiency and provides accurate location estimation for networks with 10000 sensors and beyond.

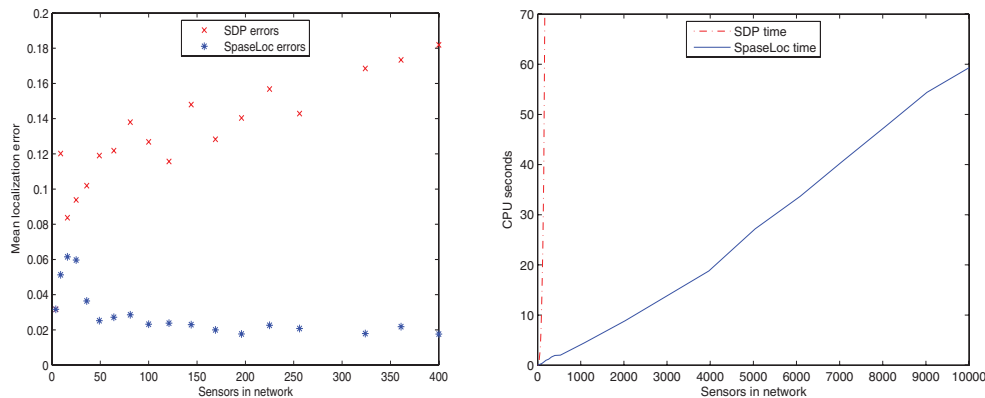


FIG. 5.1. *Accuracy and performance comparison.*

Figure 5.1 compares localization results for our SpaseLoc algorithm and the full SDP approach [2] for various sized networks. The left-hand figure shows a comparison in terms of estimation accuracy for localizing various sizes of networks when sensors are placed at the vertices of an equilateral triangle grid with 0.1 *noise_factor* added to distance measurements (data is taken from Table 4.3). It shows clearly that SpaseLoc provides much improved localization accuracy.

The right-hand graph summarizes results in terms of execution time on various network sizes. Data for the full SDP method is taken from Table 4.3, and data for SpaseLoc is taken from Table 4.4. The figure confirms near-linear complexity for SpaseLoc.

5.1. More general problems. In Jin [14], SpaseLoc is used as a building block for more general localization algorithms. A dynamic version can estimate moving sensors' locations in real time, and a three-dimensional version extends its utility further. For clustered and distributed environments, it is shown how to use SpaseLoc

in parallel (on multiple large subproblems) to obtain essentially linear complexity on clustered networks of unlimited size. Finally, a preprocessor for SpaseLoc has been developed in [14] to localize sensors in anchorless networks.

5.2. A bootstrap procedure. SpaseLoc works effectively when Step A1 (subproblem creation) finds subsensors connected to at least 3 anchors. A difficult situation arises if there are more than 3 anchors in the network but no subsensor is directly connected to 3 anchors. A network with anchors placed at the borders of the region is such an example. SpaseLoc's subproblems will involve sensors connected to only 2 or 1 anchors, leading to a less accurate final solution.

When there is sufficient connection information for sensors to be indirectly connected to at least 3 anchors through other sensors, the full SDP approach can find a solution. We are developing a procedure to choose a subproblem in the above situation. It will include the anchors, certain subsensors, and the sensors on each shortest path from a subsensor to an anchor.

5.3. Alternative subproblem solvers. At present, most of the SpaseLoc subproblems are solved by the SDP approach of Biswas and Ye [2]. This is an approximation method that may produce large errors with noisy data. A recent development by Biswas et al. [4] adds regularization terms to the SDP problem and uses a gradient-descent method to refine the SDP solution. Significant accuracy improvement is reported. An advantage of SpaseLoc is that it can solve each subproblem by any method that is effective on small networks. Our next step is to experiment with such approaches, including that of [4] and various triangulation-based methods.

Acknowledgments. We gratefully acknowledge valuable technical advice from Profs. Henry Wolkowicz, Scott Rogers, and Daniel Frances and two perceptive referees. Thanks also to Dr. Steve Benson for his expert advice on using the DSDP5.0 solver and to Prof. Kenneth Holmström for his help in fine-tuning parts of the MATLAB implementation of SpaseLoc. We are further indebted to Robert Bosch Corporation and the talented Bosch RTC team: Sharmila Ravula, Lakshmi Venkatraman, Bhaskar Srinivasan, Abtin Keshavarzian, Hauke Schmidt, and Karsten Funk; they provided the inspiration and vision for our algorithms to be practical.

REFERENCES

- [1] S. J. BENSON, Y. YE, AND X. ZHANG, *DSDP Website*, <http://www-unix.mcs.anl.gov/~benson> or <http://www.stanford.edu/~yyye/Col.html> (1998–2005).
- [2] P. BISWAS AND Y. YE, *Semidefinite programming for ad hoc wireless sensor network localization*, in Proceedings of the Third International Symposium on Information Processing in Sensor Networks, Berkeley, CA, 2004.
- [3] P. BISWAS AND Y. YE, *A distributed method for solving semidefinite programs arising from ad hoc wireless sensor network localization*, in Multiscale Optimization Methods and Applications, Nonconvex Optim. Appl. 82, Springer, New York, 2006.
- [4] P. BISWAS, T. LIANG, K. TOH, T. WANG, AND Y. YE, *Semidefinite programming approaches for sensor network localization with noisy distance measurements*, IEEE Trans. Automation Sci. Engrg., to appear.
- [5] S. BOYD, L. EL GHAOUI, E. FERON, AND V. BALAKRISHNAN, *Linear Matrix Inequalities in System and Control Theory*, SIAM Stud. Appl. Math. 15, SIAM, Philadelphia, 1994.
- [6] N. BULUSU, J. HEIDEMANN, AND D. ESTRIN, *GPS-less Low Cost Outdoor Localization for Very Small Devices*, Technical report 00-729, Computer Science Department, University of Southern California, Los Angeles, CA, 2000.
- [7] D. CULLER AND W. HONG, *Wireless sensor networks*, Comm. ACM, 47 (2004), pp. 30–33.
- [8] L. DOHERTY, L. EL GHAOUI, AND K. PISTER, *Convex position estimation in wireless sensor networks*, in Proceedings of the IEEE Infocom 2001, Anchorage, AK, 2001, pp. 1655–1663.
- [9] A. DRAGOON, *Small wonders*, CIO Magazine, January 15, 2005.

- [10] D. GANESAN, B. KRISHNAMACHARI, A. WOO, D. CULLER, D. ESTRIN, AND S. WICKER, *An Empirical Study of Epidemic Algorithms in Large-Scale Multihop Wireless Networks*, Report UCLA/CSD-TR-02-0013, Computer Science Department, University of California at Los Angeles, Los Angeles, CA, 2002.
- [11] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore, 1996.
- [12] J. HIGHTOWER AND G. BORIELLO, *Location systems for ubiquitous computing*, IEEE Computer, 34 (2001), pp. 57–66.
- [13] A. HOWARD, M. J. MATARIC, AND G. S. SUKHATME, *Relaxation on a mesh: A formalism for generalized localization*, in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS01), Maui, HI, 2001, pp. 1055–1060.
- [14] H. H. JIN, *Scalable Sensor Localization Algorithms for Wireless Sensor Networks*, Ph.D. thesis, University of Toronto, Toronto, Canada, 2005. (Joint research conducted at Stanford University.)
- [15] M. LAWLOR, *Small systems, big business*, Signal Magazine, January 2005.
- [16] MATLAB 6.5, Release 13 with Service Pack 1, The MathWorks, Inc., Natick, MA, 2003.
- [17] D. NICULESCU AND B. NATH, *Ad hoc positioning system*, in Proceedings of the IEEE GlobeCom 2001, San Antonio, TX, 2001, pp. 2926–2931.
- [18] A. RICADELA, *Sensors everywhere*, Information Week, January 24, 2005.
- [19] C. SAVARESE, J. RABAHEY, AND K. LANGENDOEN, *Robust positioning algorithm for distributed ad hoc wireless sensor networks*, in Proceedings of the USENIX Technical Annual Conference, Monterey, CA, 2002.
- [20] A. SAVVIDES, C.-C. HAN, AND M. B. SRIVASTAVA, *Dynamic fine-grained localization in ad hoc networks of sensors*, in Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MOBICON), Rome, Italy, 2001, pp. 166–179.
- [21] A. SAVVIDES, H. PARK, AND M. B. SRIVASTAVA, *The bits and flops of the n-hop multilateration primitive for node localization problems*, in Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02), Atlanta, GA, ACM Press, New York, 2002, pp. 112–121.
- [22] Y. SHANG, W. RUML, Y. ZHANG, AND M. FROMHERZ, *Localization from mere connectivity*, in Proceedings of the Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing, Annapolis, MD, ACM Press, New York, 2003, pp. 201–212.
- [23] J. F. STURM, *Let SeDuMi Seduce You*, <http://fewcal.kub.nl/sturm/software/sedumi.html> (October 2001).
- [24] R. SZEWCZYK, E. OSTERWEIL, J. POLASTRE, M. HAMILTON, A. MAINWARING, AND D. ESTRIN, *Habitat monitoring with sensor networks*, Comm. ACM, 47 (2004), pp. 34–44.
- [25] P. TSENG, *SOCP relaxation for nonconvex optimization*, presented at ICCOPT 1, Rensselaer Polytechnic Institute, Troy, NY, 2004.