Robust System Multiangulation Using Subspace Methods

Joshua N. Ash* and Lee C. Potter
Department of Electrical and Computer Engineering
Ohio State University
Columbus, OH 43210
[ashj, potter]@ece.osu.edu

ABSTRACT

Sensor location information is a prerequisite to the utility of most sensor networks. In this paper we present a robust and low-complexity algorithm to self-localize and orient sensors in a network based on angle-of-arrival (AOA) information. The proposed non-iterative subspace-based method is robust to missing and noisy measurements and works for cases when sensor orientations are either known or unknown. We show that the computational complexity of the algorithm is $\mathcal{O}(mn^2)$, where m is the number of measurements and n is the total number of sensors. Simulation results demonstrate that the error of the proposed subspace algorithm is only marginally greater than an iterative maximum-likelihood estimator (MLE), while the computational complexity is two orders of magnitude less. Additionally, the iterative MLE is prone to converge to local maxima in the likelihood function without accurate initialization. We illustrate that the proposed subspace method can be used to initialize the MLE and obtain near-Cramér-Rao performance for sensor localization. Finally, the scalability of the subspace algorithm is illustrated by demonstrating how clusters within a large network may be individually localized and then merged.

Categories and Subject Descriptors

C.2.4 [**Computer-communication networks**]: Distributed systems; C.3 [**Special-purpose and application-based systems**]: Signal processing systems

General Terms

Algorithms, Design

Keywords

Sensor networks, Localization, Calibration, Angle-of-arrival (AOA)

*The work of the first author was supported by an MIT Lincoln Laboratory Graduate Fellowship.

IPSN'07, April 25-27, 2007, Cambridge, Massachusetts, USA.

1. INTRODUCTION

Multiangulation (or simply angulation) represents a generalization of the familiar concept of triangulation and denotes the task of localizing a source point from a set of known locations and the angles from the known locations to the source. In the noiseless case, two known-location points are sufficient to localize a third point—hence the prevalence of triangulation—however, larger numbers of known-location points are beneficial when the angles contain noise. In this paper, we consider a novel technique for the *simultaneous* relative localization of all nodes in a sensor network from a set of angle-of-arrival (AOA) measurements. For this reason, we refer to the method as "system" multiangulation. No known-location points (called anchors) are required, however, they may be used to provide the relative estimate an absolute frame of reference.

Many sensors support AOA measurements as part of their sensing or communication subsystems. For example, sensors tracking people, wildlife, and vehicles are often equipped with phased microphone arrays which allow enhanced observations in the direction of the target of interest through beamforming [5, 4]. These arrays can also be used for AOA estimation via algorithms like MUSIC and ESPRIT [19]. Sensors may also be equipped with arrays of directional antennas in order to focus communication transmissions toward the intended destination and limit interference to neighboring sensors. With knowledge of the antenna beam patterns, a multi-antenna system with a single radio receiver may also be used for accurate AOA estimation of any inband RF source [1]. RF phased arrays, as used in current MIMO communication systems, will likely make their way into sensor networks and could also be used for AOA estimation, albeit with increased circuit complexity.

The aim of this paper is to develop a robust and low-complexity algorithm for network localization from AOA measurements. We illustrate that the system multiangulation problem may be interpreted as a subspace identification problem. When a complete set of inter-node *distance* measurements are available, classical multidimensional scaling (MDS) provides a robust subspace-based location estimate [3]. The Isomap algorithm generalizes MDS to incomplete measurement sets by replacing missing measurements with shortest-path distances [20]. Because time-of-arrival (TOA) and received-signal-strength (RSS) measurements may be converted to distance estimates, MDS and Isomap apply to these measurement modalities as well. Subspace methods were used in [16] for closed-form source localization using range differences and time differences. To the best of our

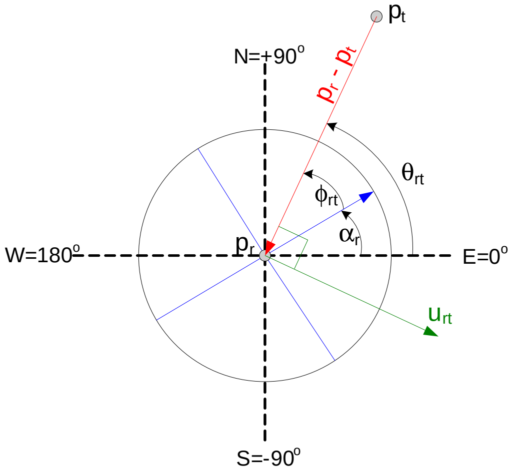

Figure 1: Global coordinate system and local coordinate system at node r

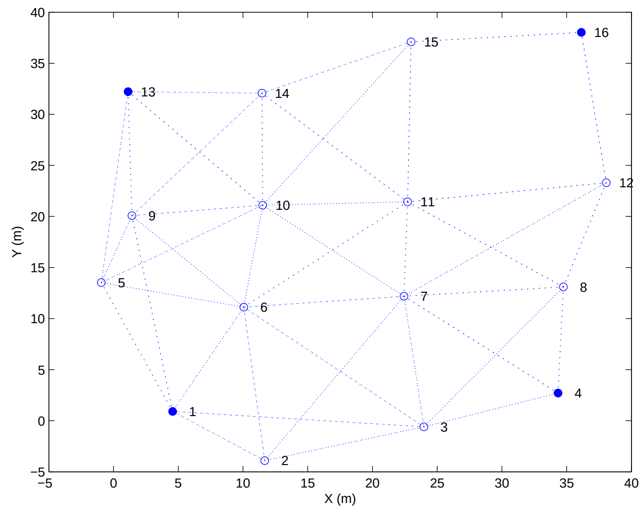

Figure 2: Sample sparse network to be localized with anchor nodes on the four corners

knowledge, subspace-based methods have not been previously applied to AOA. As such, this paper fills a gap in existing literature by providing subspace-based localization from AOA or angle-difference-of-arrival (ADOA) measurements. The algorithm is non-iterative and provides a good trade-off between complexity and performance.

Previous work on AOA-based localization considered a distributed approach where bearing-to-anchor information is propagated throughout the network allowing nodes to triangulate themselves when they obtain enough anchor bearings [11]. The strength of this approach is that it is distributed; the weakness is that bearing error accumulates as the estimates are propagated through the network. In [7], a centralized approach was taken where semidefinite programming was employed to estimate sensor locations from the intersection of AOA-derived constraint sets. Maximum-likelihood estimates of sensor positions from AOA measurements were considered in [10].

Notation: Throughout this paper capital letters represent matrices and bold lowercase letters denote column vectors. $(\cdot)^T$ and $(\cdot)^*$ represent matrix transpose and conjugate-transpose respectively, while $\mathcal{R}(A)$ and $\mathcal{N}(A)$ represent the range space and null space of matrix A, respectively.

2. SYSTEM MULTIANGULATION

The localization problem is to estimate sensor node positions $\boldsymbol{p}_i = [x_i \ y_i]^T$, $i \in \{1, ..., n\}$ from a partial set of angle-of-arrival measurements $\{\phi_{ij}\}$. In order to obtain absolute position estimates, a small subset $\mathcal{K} \subseteq \{1, ..., n\}$ of the nodes, called anchors, are assumed to have known positions and are used to resolve ambiguities in translation, rotation, and scale of a relative solution. As illustrated in Figure 1, each node r makes AOA measurements in a local coordinate system which is offset by an angle α_r from a global reference. In the global coordinate system, the AOA at receiving node r, of a transmission from node t is

$$\theta_{rt} = \phi_{rt} + \alpha_r, \tag{1}$$

where ϕ_{rt} is the measurement in r's local system. We first consider localization when the orientation angles $\{\alpha_r\}$ are known and then address the more difficult unknown orientation case.

2.1 Known orientations

When the sensor orientations $\{\alpha_r\}$ are known, we can immediately map local measurements $\{\phi_{rt}\}$ into their global frame counterparts $\{\theta_{rt}\}$ via (1). For each angle θ_{rt}, we form a unit vector

$$\boldsymbol{u}_{rt} = \begin{bmatrix} \sin\theta_{rt} \\ -\cos\theta_{rt} \end{bmatrix} \tag{2}$$

which, as illustrated in Figure 1, is orthogonal to the difference of the position vectors

$$\boldsymbol{u}_{rt}^T(\boldsymbol{p}_t - \boldsymbol{p}_r) = 0. \tag{3}$$

This construction allows us to determine a system of equations which can be solved for the node positions.

Let $\mathcal{M}$ denote the set of ordered measurement pairs; that is, $(r,t) \in \mathcal{M}$ if node r makes an AOA measurement of a transmission from t, and let $\mathcal{M}_r(i)$ and $\mathcal{M}_t(i)$ denote the receiver and transmitter of the ith element of $\mathcal{M}$, respectively. Let $U = \{U_{ij}\}$ be an $|\mathcal{M}| \times |\mathcal{M}|$ block diagonal matrix, where each block is an element of $\mathbb{R}^{2\times1}$ and $|\mathcal{M}|$ is the total number of measurements. We populate the diagonal (block) entries of U with the previously defined unit vectors as $U_{ii} = \boldsymbol{u}_{r't'}$, where $r' = \mathcal{M}_r(i)$, $t' = \mathcal{M}_t(i)$. Let $K = \{K_{ij}\}$ be an $|\mathcal{M}| \times n$ block matrix with common block size 2×2. The non-zero elements of the ith row of K are populated as $K_{it'} = I_2$ and $K_{ir'} = -I_2$, where I_2 is the 2×2 identity matrix and again $r' = \mathcal{M}_r(i)$, $t' = \mathcal{M}_t(i)$. If we let $\boldsymbol{p} = [\boldsymbol{p}_1^T \, \boldsymbol{p}_2^T \, \ldots \, \boldsymbol{p}_n^T]^T$ be the stacking of the n position vectors, then from (3) we can form the homogeneous linear system

$$U^T K \boldsymbol{p} = 0. \tag{4}$$

The function of the matrix K is to form the necessary position vector differences corresponding to each AOA measurement. An example of a three-node system with four measurements may look like

$$\underbrace{\begin{bmatrix} \boldsymbol{u}_{21} & 0 & 0 & 0 \\ 0 & \boldsymbol{u}_{31} & 0 & 0 \\ 0 & 0 & \boldsymbol{u}_{32} & 0 \\ 0 & 0 & 0 & \boldsymbol{u}_{13} \end{bmatrix}^T}_{U^T,\ |\mathcal{M}|\times 2|\mathcal{M}|} \underbrace{\begin{bmatrix} I_2 & -I_2 & 0_2 \\ I_2 & 0_2 & -I_2 \\ 0_2 & I_2 & -I_2 \\ -I_2 & 0_2 & I_2 \end{bmatrix}}_{K,\ 2|\mathcal{M}|\times 2n} \underbrace{\begin{bmatrix} \boldsymbol{p}_1 \\ \boldsymbol{p}_2 \\ \boldsymbol{p}_3 \end{bmatrix}}_{p,\ 2n\times 1} = \underbrace{\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}}_{|\mathcal{M}|\times 1} . \tag{5}$$

The relative localization solution sought is a particular solution to the linear homogeneous system (4) above. In general, the dimension of the nullspace of $U^T K$ is three. One dimension corresponds to scalings of the point configuration $\boldsymbol{p}$, another corresponds to translations in the x-direction, and the third corresponds to translations along the y-direction; that is $\mathcal{N}(U^T K) = \mathcal{R}([\boldsymbol{p}\ \boldsymbol{v}_x\ \boldsymbol{v}_y])$, where $\boldsymbol{v}_x = [1\ 0\ 1\ 0 \ldots]^T$ and $\boldsymbol{v}_y = [0\ 1\ 0\ 1 \ldots]^T$. The vectors $\boldsymbol{v}_x$ and $\boldsymbol{v}_y$ come directly from the matrix K, ($K\boldsymbol{v}_x = K\boldsymbol{v}_y = 0$) and reflect the fact that AOA measurements only depend on the relative point configuration and are invariant to scene translations. The third basis vector is, by construction, the desired point configuration $\boldsymbol{p}$ up to scale and translation. If there are an insufficient number of measurements, the dimension of $\mathcal{N}(U^T K)$ will be greater than three.

Because the translation vectors $\boldsymbol{v}_x$ and $\boldsymbol{v}_y$ do not depend on the coordinates being estimated, we can reduce the dimension of the null space by augmenting the rows of $U^T K$ with $\boldsymbol{v}_x^T$ and $\boldsymbol{v}_y^T$. Let

$$A = \begin{bmatrix} U^T K \\ \boldsymbol{v}_x^T \\ \boldsymbol{v}_y^T \end{bmatrix} \in \mathbb{R}^{(|\mathcal{M}|+2)\times(2n)}. \tag{6}$$

Then, $\mathcal{N}(A) = \mathrm{span}(\boldsymbol{p})$, and we can solve

$$A\boldsymbol{p} = 0 \tag{7}$$

using the singular value decomposition $A = U_A \Sigma_A V_A^T$. The (unit-norm) minimizing solution $\boldsymbol{p}$ of $\|A\boldsymbol{p}\|$ is

$$\hat{\boldsymbol{p}} = V_A^{(2n)}, \tag{8}$$

where $V_A^{(2n)}$ is the right singular vector corresponding to the minimum singular value of A. The point configuration so obtained will be centered at the origin and, in the noiseless case, will be a equal to the generating configuration up to translation and scale. The final estimate of $\boldsymbol{p}$ is obtained by scaling and translating $\hat{\boldsymbol{p}}$ for maximal agreement with prior information. We use the Procrustes algorithm for this purpose as described in the next subsection.

2.1.1 Transforming from relative to absolute point configurations

In order to transform our relative solution obtained above into an absolute solution, we we find a suitable transformation mapping the estimated values of the anchor nodes $\mathcal{K}$ to their *a priori* known positions.

Let the $|\mathcal{K}| \times 2$ matrix $P = [\boldsymbol{p}_{\mathcal{K}(1)}, \boldsymbol{p}_{\mathcal{K}(2)}, \ldots]^T$ denote the matrix of *a priori* known positions, and similarly let $\hat{P} = [\hat{\boldsymbol{p}}_{\mathcal{K}(1)}, \hat{\boldsymbol{p}}_{\mathcal{K}(2)}, \ldots]^T$ form the previously obtained estimates of this subset of nodes. We use the Procrustes algorithm to find the squared-error minimizing translation, rotation, and scale of the points $\hat{P}$ into P. That is, we seek a scale factor s, an orthogonal rotation matrix R, and a translation vector $\boldsymbol{t} = [\boldsymbol{t}_x\ \boldsymbol{t}_y]^T$ that minimizes

$$\left\| P - (s\hat{P}R + \mathbf{1}\boldsymbol{t}^T) \right\|_F, \tag{9}$$

where the norm is in the Frobenius sense and $\mathbf{1}$ is a vector of all ones. We then apply these transformations to all of the estimated points in order to obtain the complete scene estimate.

The solution to this problem is a slight generalization of the orthogonal Procrustes problem and is given in [15]. Let $J = (I - |\mathcal{K}|^{-1}\mathbf{1}\mathbf{1}^T)$, and compute the SVD $U_p \Sigma_p V_p^T = P^T J \hat{P}$. The desired transformation parameters are

$$R = V_p U_p^T \tag{10}$$

$$s = (tr\ P^T J \hat{P} R)/(tr\ \hat{P}^T J \hat{P}) \tag{11}$$

$$\boldsymbol{t} = |\mathcal{K}|^{-1}(P - s\hat{P}R)^T \mathbf{1}. \tag{12}$$

For the known orientation case, the scene rotation is not ambiguous. That is, the rotation R above will always be identity (or a reflection matrix), however, rotation will be needed below for the case of unknown orientations. The complexity of determining rotation, translation, and scale using the Procrustes method is very low because the SVD only involves a 2×2 matrix, for any number of anchors and total sensors. The complexity of obtaining the relative estimate $\hat{\boldsymbol{p}}$ is addressed in Section 4 after considering the unknown orientation case.

2.2 Unknown orientations

Unknown orientations $\{\alpha_r\}$ are nuisance parameters in the localization problem, however their estimates are useful quantities in many sensor applications. In our approach, we first estimate the orientations up to a common offset and then apply the point estimation procedure above for known orientations. We begin by describing the procedure for the case of complete measurements where all pairs of sensors make AOA measurements in both directions. This is then extended to the partial measurement case.

2.2.1 Complete measurements

Using the fact that $\theta_{rt} = \theta_{tr} + \pi$ we have

$$\alpha_r - \alpha_t = \phi_{tr} - \phi_{rt} + \pi \qquad \text{for all } r \neq t. \tag{13}$$

In order to simplify the modulo 2π arithmetic involved in angular calculations, we represent the sensor orientation angles as the phase of points in the complex plane, $\boldsymbol{a}(\boldsymbol{\alpha}) = [e^{i\alpha_1},\ \ldots,\ e^{i\alpha_n}]^T$, where $\boldsymbol{\alpha} = [\alpha_1,\ \ldots,\ \alpha_n]^T$ is the vector of unknown orientations. Let $\Phi = \{\phi_{rt}\}$ denote the matrix of AOA measurements, from which we form

$$\Psi = \Phi^T - \Phi + \pi \mathbf{1}_n \mathbf{1}_n^T, \quad \text{and} \tag{14}$$

$$B = e^{i\Psi} + 2\boldsymbol{I}, \tag{15}$$

where $e^{i\Psi}$ denotes elementwise exponentiation of Ψ, not the conventional matrix exponential.

We now have the fundamental relation

$$\boldsymbol{a}(\boldsymbol{\alpha})\ \boldsymbol{a}(\boldsymbol{\alpha})^* = B \tag{16}$$

which holds in the noiseless case. Let $V_B \Lambda_B V_B^*$ denote the eigen-decomposition of B. The closest rank-one approximation (in the Frobenius norm sense) to B is $\lambda_n \boldsymbol{v}_n \boldsymbol{v}_n^*$, where

λ_n is the largest eigenvalue of B and $\boldsymbol{v}_n$ is the associated eigenvector. We may then estimate $\boldsymbol{a}(\boldsymbol{\alpha})$ as

$$\hat{\boldsymbol{a}} = \boldsymbol{v}_n \tag{17}$$

and compute the orientation estimates as the phase of each complex element of $\hat{\boldsymbol{a}}$

$$\hat{\alpha}_i = \angle \hat{a}_i. \tag{18}$$

In the noiseless case, the orientation angles so obtained will equal the true angles up to a constant, that is $\boldsymbol{\alpha} = \hat{\boldsymbol{\alpha}} + \alpha_0 \mathbf{1}$. The ambiguity of α_0 is due to the fact that the elements of B only depend on the differences between elements of $\boldsymbol{\alpha}$ as seen in (13) and (15). The scalar α_0 represents a common orientation offset which results in an overall unknown rotation of the relative scene estimate. This last fact is seen by noting that if $\boldsymbol{p}_r$ and $\boldsymbol{p}_t$ satisfy (3), then $\tilde{p}_r = R_{\alpha_0} \boldsymbol{p}_r$ and $\tilde{p}_t = R_{\alpha_0} \boldsymbol{p}_t$ satisfy

$$\begin{bmatrix} \sin(\theta_{rt} + \alpha_0) \\ -\cos(\theta_{rt} + \alpha_0) \end{bmatrix}^T (\tilde{p}_r - \tilde{p}_t) = 0, \tag{19}$$

where

$$R_{\alpha_0} = \begin{bmatrix} \cos\alpha_0 & -\sin\alpha_0 \\ \sin\alpha_0 & \cos\alpha_0 \end{bmatrix} \tag{20}$$

is a counterclockwise rotation matrix by α_0.

Thus, the unknown α_0 corresponds to an overall unknown scene rotation which can be disambiguated with prior knowledge – along with scale and translation. Therefore, we can use the $\hat{\boldsymbol{\alpha}}$ estimates (18) to determine the unit vectors $\{\boldsymbol{u}_{rt}\}$ and proceed with localization as described earlier for the known-orientation case.

2.2.2 *Partial measurements*

The problem is complicated slightly in practice where finite communication ranges makes it unrealistic to assume that each node obtains AOA measurements from all other nodes – that is, some elements of Φ, and consequently B, may be unknown.

An unknown element $B_{rt} = e^{i(\alpha_r - \alpha_t)}$ may be derived from any known path from r to t. Consider a path

$$\mathcal{P} = \{(r, n_1), (n_1, n_2), ...(n_k, t)\} \tag{21}$$

with k intermediate nodes. We then have that

$$\begin{aligned} B_{rt} &= e^{i(\alpha_r - \alpha_t)} \\ &= e^{i(\alpha_r - \alpha_{n_1})} \cdot e^{i(\alpha_{n_1} - \alpha_{n_2})} \cdot \ldots \cdot e^{i(\alpha_{n_k} - \alpha_t)} \\ &= \prod_{(m,n) \in \mathcal{P}} B_{mn} \end{aligned} \tag{22}$$

and see that B_{rt} may be estimated as the product of entries along any path from r to t for which measurements are known. In the case of noisy measurements a minimum-length path should be used in order to decrease the variance of the resulting estimate. Conventional shortest-path algorithms, such as Floyd-Warshall and Dijkstra, are useful here. As an extension to this technique, it may be desirable to perform a *weighted* rank-one approximation of B which minimizes the weighted Frobenius norm

$$J(a) = \sum_{(r,t)} W_{rt}((aa^*)_{rt} - B_{rt})^2, \tag{23}$$

where the weights $W_{rt} = 1/h_{rt}$ and h_{rt} is equal to the number of hops used in establishing entry B_{rt}. Unfortunately, the known methods for computing weighted low-rank approximations are all iterative [17], unlike the unweighted case which is determined via a single eigen-decomposition. Although not considered here, missing entries in B could also be constructed from a weighted average among multiple paths. Our implementation is based on the shortest path using Dijkstra's algorithm and the (unweighted) eigen-approximation (17).

The entire algorithm, termed Robust Angulation using Subspace Techniques (RAST), is summarized in Table 1. In the noiseless case, RAST yields the exact solution. Performance in noise is considered in the next section.

Table 1: Robust Angulation using Subspace Techniques (RAST), algorithm summary

- if sensor orientations, α, are unknown
  - calculate B using (15)
  - if B has missing entries from incomplete measurements
    - use shortest path construction (22) to complete entries
  - compute eigen-decomposition $B = V_B \Lambda_B V_B^*$
  - estimate orientations as the phase of the elements of the eigenvector corresponding to the largest eigenvalue.
- calculate global-frame arrival angles $\{\theta_{rt}\}$ from (1) and compute unit vectors as in (2)
- form matrix A as in (6) and compute its SVD $A = U_A \Sigma_A V_A^T$
- estimate the relative point configuration from the rightmost singular vector of A, $\hat{\boldsymbol{p}} = V_A^{(2n)}$
- if absolute positioning (scale, translation, rotation) are required, use Procrustes equations (10), (11), (12) to estimate these parameters from $\hat{\boldsymbol{p}}$ and anchor points

3. PERFORMANCE IN NOISE AND THE CRB

Through simulation, we evaluate the performance of our algorithm on the sparse network depicted in Figure 2 and compare the localization performance to the Cramér-Rao bound (CRB) and the maximum-likelihood estimator (MLE). In the figure, the measurement radius has been limited to 20 m making the network relatively sparse. The measurements $\{\tilde{\phi}_{rt}\}$ are Gaussian perturbations of the true values: $\tilde{\phi}_{rt} = \phi_{rt} + n_{rt}$, $n_{rt} \sim N(0, \sigma_\phi^2)$, where σ_ϕ^2 is the common observation variance. The four corner nodes of the network serve as anchors.

Let $\boldsymbol{\Theta} = \{x_i \; y_i \; \alpha_i\}_{i=1:n}$ denote the parameter vector to be estimated, and let $\phi(\boldsymbol{\Theta})$ be the resulting AOA measurements that would result from a proposed value of $\boldsymbol{\Theta}$. The MLE for this problem is equal to the (non-linear) least squares estimate [10]

$$\hat{\boldsymbol{\Theta}}_{mle} = \arg\min_{\boldsymbol{\Theta}} J(\boldsymbol{\Theta}), \tag{24}$$

where $J(\boldsymbol{\Theta})$ is the cost function

$$J(\boldsymbol{\Theta}) = \sum_{(r,t) \in \mathcal{M}} \left(\tilde{\phi}_{rt} - \phi(\boldsymbol{\Theta})_{rt} \right)^2. \tag{25}$$

As a performance metric we consider the scene RMS error defined as

$$E_{rms} = \left(\frac{1}{n - |\mathcal{K}|} \sum_{i \notin \mathcal{K}} E[\hat{d}_i^2] \right)^{\frac{1}{2}}, \qquad (26)$$

where $E[\hat{d}_i^2]$ is the expected value of the squared distance between an estimate and the true position of node i. For the MLE and RAST estimators, this expected quantity is derived from a large number of simulations. A lower bound on the minimum possible error achievable by any unbiased estimator is given by the Cramér-Rao bound (CRB) which is derived from the inverse of the Fisher information matrix (FIM) [13]. The FIM for this problem is derived in [1, 12].

3.1 Simulations

In Figure 3 we plot the localization performance of the RAST algorithm as a function of the AOA measurement error σ_ϕ for the network in Figure 2—based on 1000 Monte Carlo measurement realizations. We consider a wide range in the quality of AOA estimates available; with errors ranging from good ($\sigma_\phi = 2.5°$) to very poor ($\sigma_\phi = 20°$). As a point of reference, AOA errors of approximately $3° - 5°$ were experimentally observed for an RF-based system in [1] and for an ultrasound-based system in [14]. From Figure 3, we see that the estimation error is approximately 60% greater than the minimum possible CRB-predicted error when $\sigma_\phi = 20°$ and sensor orientations are known. When the measurement noise decreases to $\sigma_\phi = 2.5°$, the estimation error exceedes the CRB by only 21%. Similarly, for the unknown-orientation case, the estimation error ranges from 29% to 55% above the CRB as σ_ϕ varies from 2.5° to 20°.

Figures 4 and 5 present simulation results comparing RAST performance with maximum-likelihood performance in localizing the sample network of Figure 2. The MLE was obtained as an iterative solution to the nonlinear least-squares problem (24) using Matlab's implementation of the interior-reflective Newton method [6]. In the unknown orientation case, MLE convergence is very sensitive to the initial value given to the algorithm and is prone to converge to local maxima. These figures consider three different initialization mechanisms for the MLE: *(i)* the true point positions (used as a benchmark), *(ii)* random initial points drawn uniformly from the 40m × 40m scene, and *(iii)* the output from our subspace estimator.

When the sensor orientations are known, we see from Figure 4 that all of the ML estimates have nearly identical error and that they outperform RAST errorwise. However, as illustrated in Table 2, the closed-form nature of RAST results in a runtime that is nearly 2 orders of magnitude less than the iterative MLE. Algorithm runtimes were measured with Matlab's `cputime` function.

When orientations are unknown, the cost function (25) becomes plagued with local maxima and the MLE optimization has a very difficult time converging to the global maximum. In the case of random initialization the MLE always converged to distant local maxima resulting in errors significantly outside the plot region of Figure 5. However, the RAST algorithm provided estimates sufficiently close to the attraction region of the optimal solution that MLE initialized from RAST was nearly equivalent to initialization with true values. As such, RAST-initialized maximum-likelihood estimates nearly achieve CRB performance. For $2.5° \leq \sigma_\phi \leq 12°$, the amount that these estimates exceed the CRB by ranges from only 1% to 4%, for both known and unknown orientations. Runtimes for the unknown orientation case are given in Table 3.

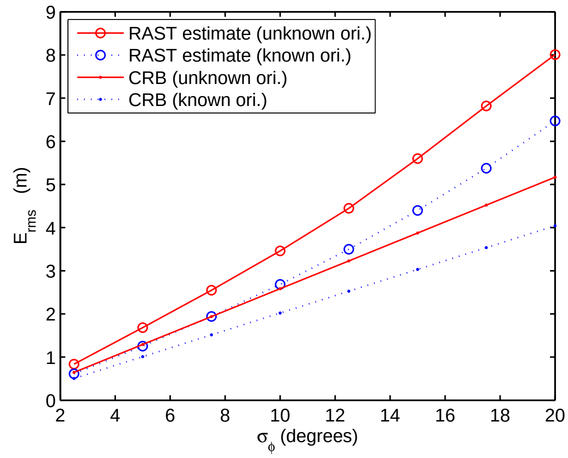

Figure 3: Performance of the RAST localization method for known and unknown sensor orientations compared to the Cramér-Rao lower bound

Table 2: *Orientations known* – Performance and CPU time characteristics of each estimator. $\sigma_\phi = 5°$

Estimator	RMS Error (m)	CPU time per estimate (ms)
RAST	1.23	9
ML(true)	1.04	384
ML(rast)	1.04	377
ML(rand)	1.04	823
CRB	1.01	*n.a.*

4. COMPLEXITY

The dimensions of the matrix A in (6) are $(|\mathcal{M}|+2) \times 2n$, and for practical cases the total number of measurements $|\mathcal{M}| > 2n$, making the complexity of evaluating the SVD $\mathcal{O}\left((|\mathcal{M}|+2)(2n)^2\right)$. For unknown orientations, the matrix B in (15) is $n \times n$ and its eigen-decomposition $\mathcal{O}(n^3)$. Using the Dijkstra or Floyd-Warshall algorithms for shortest path construction to fill in missing entries in the matrix B has worst-case complexity $\mathcal{O}(n^3)$ [2, Ch. 5]. Thus, the overall

Table 3: *Orientations unknown* – Performance and CPU time characteristics of each estimator. $\sigma_\phi = 5°$

Estimator	RMS Error (m)	CPU time per estimate (ms)
RAST	1.60	48
ML(true)	1.30	670
ML(rast)	1.30	662
ML(rand)	3418	19258
CRB	1.29	*n.a.*

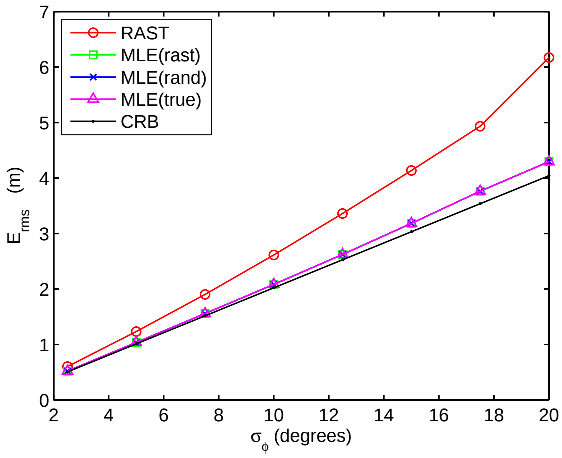

Figure 4: *Orientations known* – Comparison of the RAST estimator to the iterative maximum-likelihood estimator initialized in different ways. All maximum-likelihood estimates are nearly identical for this case.

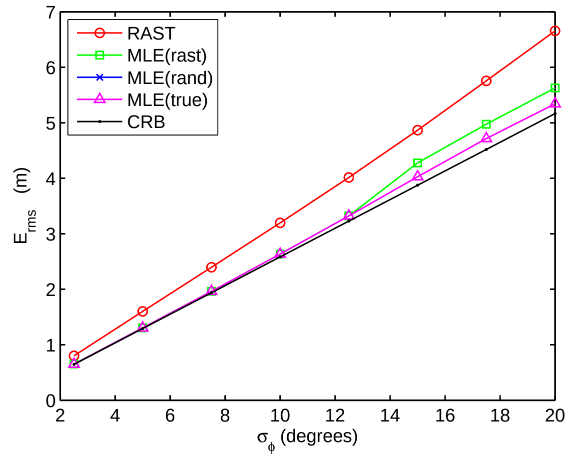

Figure 5: *Orientations unknown* – For this case, random initialize of the MLE failed to achieve the global maximum and errors are significantly outside of the plot region. MLE initialization from RAST produces results nearly as good as true-value initialization (used as a benchmark).

complexity of the algorithm is $\mathcal{O}(|\mathcal{M}|\, n^2)$ for both known and unknown orientation angles.

There are a number of numerical linear algebra techniques that can reduce this complexity. The matrix A is sparse and we need only the least dominant right singular vector, not the entire SVD; thus using Lanczos method reduces the storage and compute time required[8, Ch. 9]. Using Lanczos method and the RAST algorithm, we observe that a 2 GHz desktop can simultaneously localize a 2000 node network in approximately 15 sec when each node makes 10 AOA measurements to its neighbors.

If new AOA measurements are subsequently made after an initial localization solution obtained, the matrix A can be augmented with new rows and the original nullspace basis $\boldsymbol{p}$ can be *updated* in $\mathcal{O}(n^2)$ operations without having to compute the SVD of the augmented matrix from scratch [8,§12.5.5; 18]. The situation of a new sensor node joining a previously calibrated network is more complicated as both new rows and columns must be appended to A. In this case, a previously computed SVD of A can be updated in two steps as row augmentations to A and A^T as described in [9]. This update can be performed in $\mathcal{O}\left((|\mathcal{M}|+n)n\right)$ operations.

5. MERGING SUBGRAPHS

The subspace-based algorithm presented in this paper is centralized, however, it naturally accommodates the merging of localization solutions from multiple subgraphs into one complete solution. This fits the popular hierarchical model for large sensor networks where the entire network is divided into clusters, and each cluster is managed by a cluster-head. Using the methods described in this section, each cluster may perform localization of its constituent nodes using the previously described algorithm. All of the clusters may then be "stitched" together into a complete network. For clarity of presentation, we describe the procedure for a 2-cluster system, as illustrated in Figure 6.

5.1 Known orientations

Let $\boldsymbol{p}^{(1)}$ denote a relative localization solution for cluster one. That is, $U_1 K_1 \boldsymbol{p}^{(1)} = 0$, where U_1 is derived solely from measurements between nodes in cluster one, and K_1 forms the necessary differences between positions of nodes in cluster one. Similarly, let the solution for the second cluster be $U_2 K_2 \boldsymbol{p}^{(2)} = 0$. Without loss of generality, we will assume the position of $\boldsymbol{p}^{(1)}$ to be fixed and determine the transformation of $\boldsymbol{p}^{(2)}$ necessary to merge the two solutions in a manner that is consistent with the measurements made *between* the two clusters. When sensor orientations are known, this amounts to determining a suitable translation and scaling of $\boldsymbol{p}^{(2)}$. In particular, we seek

$$\tilde{\boldsymbol{p}}^{(2)} = s\,\boldsymbol{p}^{(2)} + x\,\boldsymbol{v}_x + y\,\boldsymbol{v}_y = Q\,\boldsymbol{\delta}, \tag{27}$$

where $Q = [\boldsymbol{p}^{(2)}\ \boldsymbol{v}_x\ \boldsymbol{v}_y]$, and the elements of $\boldsymbol{\delta} = [s\ x\ y]^T$ represent scale, x-translation, and y-translation, respectively.

The linear system (4) for the entire network may be written

$$\begin{bmatrix} U_1^T & 0 & 0 \\ 0 & U_x^T & 0 \\ 0 & 0 & U_2^T \end{bmatrix} \begin{bmatrix} K_1 & 0 \\ K_x^{(1)} & K_x^{(2)} \\ 0 & K_2 \end{bmatrix} \begin{bmatrix} \boldsymbol{p}^{(1)} \\ \boldsymbol{p}^{(2)} \end{bmatrix} = 0, \tag{28}$$

where U_x is formed from cross-cluster measurements (transmitter in one cluster, receiver in another) and $[K_x^{(1)} K_x^{(2)}]$ forms the necessary vector differences between cross-cluster pairs. Rewriting (28) as

$$\begin{bmatrix} U_1^T K_1 & 0 \\ U_x^T K_x^{(1)} & U_x^T K_x^{(2)} \\ 0 & U_2^T K_2 \end{bmatrix} \begin{bmatrix} \boldsymbol{p}^{(1)} \\ \boldsymbol{p}^{(2)} \end{bmatrix} = 0, \tag{29}$$

we see that the total network gives three systems of equations (i) $U_1^T K_1 \boldsymbol{p}^{(1)} = 0$, (ii) $U_x^T K_x^{(1)} \boldsymbol{p}^{(1)} + U_x^T K_x^{(2)} \boldsymbol{p}^{(2)} = 0$, and (iii) $U_2^T K_2 \boldsymbol{p}^{(2)} = 0$. Equations (i) and (iii) are solved

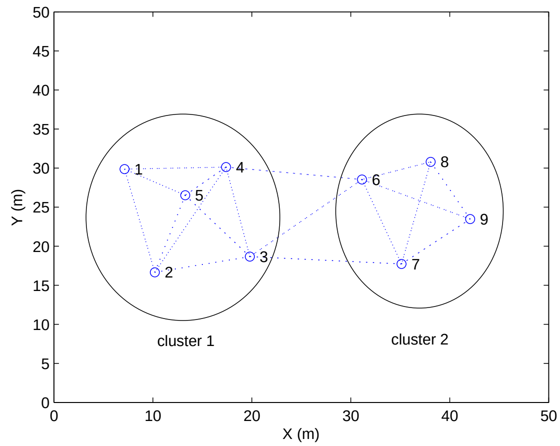

Figure 6: A large network can be localized in clusters which may subsequently be merged together according to cross-cluster measurements.

by performing relative localization within each cluster (as already assumed), and equation (ii) provides the constraints necessary to determine how the clusters themselves are positioned relative to one another. Combining (ii) and (27) we can estimate the translation and scale parameters as

$$\hat{\boldsymbol{\delta}} = -(U_x^T K_x^{(2)} Q)^\dagger U_x^T K_x^{(1)} \boldsymbol{p}^{(1)}, \tag{30}$$

where $(\cdot)^\dagger$ denotes pseudo-inverse.

5.2 Unknown orientations

When the orientations of the sensors $\{\alpha_i\}$ are unknown, the relative cluster estimates $\boldsymbol{p}^{(1)}$ and $\boldsymbol{p}^{(2)}$ have arbitrary orientations. Thus, in addition to translation and scale, we also seek a rotation angle Δ of $\boldsymbol{p}^{(2)}$.

Let $\hat{\boldsymbol{a}}^{(i)}$ denote the complex orientation estimates from the relative localization of cluster i (see (17)), and let $\tilde{\boldsymbol{a}}^{(i)} \in \mathbb{C}^{n^{(i)}}$ denote the subset of $\hat{\boldsymbol{a}}^{(i)}$ corresponding to the $n^{(i)}$ sensors of cluster i involved in cross-cluster measurements. For the example in Figure 6, $n^{(1)} = 2$ and $\tilde{\boldsymbol{a}}^{(1)} = [e^{i\alpha_3}\ e^{i\alpha_4}]^T$. Form the $(n^{(1)}+n^{(2)}) \times (n^{(1)}+n^{(2)})$ matrix B_x as in (15) using the cross-cluster measurements and compute its rank-1 approximation using eigen-decomposition as in (17)

$$\hat{\boldsymbol{a}}_x\, \hat{\boldsymbol{a}}_x^* \approx B_x. \tag{31}$$

As before, the elements of $\hat{\boldsymbol{a}}_x$ are partitioned as $\hat{\boldsymbol{a}}_x^{(1)}$ and $\hat{\boldsymbol{a}}_x^{(2)}$, representing the elements from cluster 1 and cluster 2, respectively. Finally, rotational offsets $\Delta^{(i)}$ need to be estimated for each cluster in order for the previous orientation estimates to coincide with those of the cross-cluster measurements. Specifically, we solve for real scalars $\Delta^{(1)}$ and $\Delta^{(2)}$ such that

$$\left\| \begin{bmatrix} e^{i\Delta^{(1)}}\ \tilde{\boldsymbol{a}}^{(1)} \\ e^{i\Delta^{(2)}}\ \tilde{\boldsymbol{a}}^{(2)} \end{bmatrix} - \begin{bmatrix} \hat{\boldsymbol{a}}_x^{(1)} \\ \hat{\boldsymbol{a}}_x^{(2)} \end{bmatrix} \right\|_2^2 \tag{32}$$

is minimized:

$$\Delta^{(i)} = \angle\left((\tilde{\boldsymbol{a}}^{(i)*} \hat{\boldsymbol{a}}_x^{(i)}) / (\tilde{\boldsymbol{a}}^{(i)*} \tilde{\boldsymbol{a}}^{(i)}) \right), \quad i = 1, 2. \tag{33}$$

Since, as shown in Section 2.2.1, uniformly adjusting the orientation angles by Δ corresponds to a rotation of estimates by Δ, we can apply a rotation of $\Delta = \Delta^{(2)} - \Delta^{(1)}$ to the relative cluster estimate $\boldsymbol{p}^{(2)}$. Translation and scale may then be determined as in the previous section.

Finally, we note that the process of aligning two clusters only utilizes cross-cluster measurements and relative cluster-position estimates for sensors involved in these measurements[1]. Therefore, the transmission and computation requirements for merging are proportional to the number of cross-cluster measurements—which is expected to be much smaller than the number of intra-cluster measurements.

6. CONCLUSIONS

In this paper we presented a new localization algorithm that provides both sensor position and orientation estimates from AOA measurements. The algorithm is subspace based in that it relies on singular-value and eigen-decompositions to identify subspaces associated with position and orientation vectors. The algorithm compared favorably to the Cramér-Rao bound on estimation error variance and was shown to have complexity $\mathcal{O}(|\mathcal{M}|\, n^2)$, where $|\mathcal{M}|$ is the number of measurements and n is the total number of sensors. In Monte Carlo simulations, the closed-form nature of the algorithm resulted in a runtime nearly 2 orders of magnitude less than an iterative maximum-likelihood estimator. The fast and robust nature of the algorithm makes it well-suited to initialize high fidelity routines, such as the MLE. Finally, while the algorithm is centralized, it was demonstrated how disjoint subnetwork clusters could be localized independently and subsequently merged together with minimal overhead.

7. REFERENCES

[1] J. Ash and L. Potter. Sensor network localization via received signal strength measurements with directional antennas. *Proc. 42nd Annual Allerton Conference on Communication, Control, and Computing*, pages 1861–1870, Monticello, IL, Sep 2004.

[2] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, 2nd edition, 1992.

[3] I. Borg and P. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer, 1997.

[4] C. Chen, A. M. Ali, and H. Wang. Design and testing of robust acoustic arrays for localization and enhancement of several bird sources. *Proc. Information Processing in Sensor Networks*, pages 268–275, 2006.

[5] J. C. Chen, K. Yao, and R. E. Hudson. Source localization and beamforming. *IEEE Signal Processing Magazine*, 19(2):30–39, 2002.

[6] T. F. Coleman and Y. Li. An interior, trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6:418–445, 1996.

[7] L. Doherty, L. El Ghaoui, and K. Pister. Convex position estimation in wireless sensor networks. *Proc. INFOCOM*, 3:1655–1663, 2001.

[1]From (30), $\hat{\boldsymbol{\delta}}$ appears to depend on $\boldsymbol{p}^{(1)}$ and $\boldsymbol{p}^{(2)}$, however $K_x^{(i)}$ only has non-zero columns corresponding to sensors involved in cross-cluster measurements.

[8] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 3^{rd} edition, 1996.

[9] M. Gu and S. C. Eisenstat. A stable and fast algorithm for updating the singular value decomposition. Technical Report YALEU/DCS/RR-966, New Haven, CT, 1993.

[10] R. L. Moses, D. Krishnamurthy, and R. Patterson. An auto-calibration method for unattended ground sensors. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 3(4):III–2941–III–2944, May 13-17, 2002.

[11] D. Niculescu and B. Nath. Ad hoc positioning system (APS) using AOA. *Proc. INFOCOM*, 3:1734 – 1743, 2003.

[12] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero III, R. L. Moses, and N. S. Correal. Locating the nodes: cooperative localization in wireless sensor networks. *IEEE Signal Processing Magazine*, 22(4):54–69, July 2005.

[13] H. V. Poor. *An Introduction to Signal Detection and Estimation*. Springer, New York, 2^{nd} edition, 1994.

[14] N. Priyantha, A. Miu, H. Balakrishnan, and S. Teller. The cricket compass for context-aware mobile applications. *Proc. of the 6th ACM MOBICOM*, 2001.

[15] P. H. Schönemann and R. M. Carroll. Fitting one matrix to another under choice of a central dilation and a rigid motion. *Psychometrika*, 35:245 – 255, 1970.

[16] J. Smith and J. Abel. Closed-form least-squares source location estimation from range-difference measurements. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-35(12), 1987.

[17] N. Srebro and T. Jaakkola. Weighted low-rank approximations. *Proc. of the* 20^{th} *International Conf. on Machine Learning*, 2003.

[18] G. Stewart. An updating algorithm for subspace tracking. *IEEE Trans. on Signal Processing*, 40(6):1535–1541, June 1992.

[19] P. Stoica and R. Moses. *Spectral Analysis of Signals*. Prentice Hall, 2005.

[20] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.