In: Unattended Ground Sensor Technologies and Applications IV (Proceedings of SPIE Vol. 4743), Orlando, FL, 2-5 April 2002.

Self-Calibration of Sensor Networks

Randolph L. Moses and Robert Patterson

Department of Electrical Engineering, The Ohio State University, Columbus, OH 43210 USA

ABSTRACT

We consider the problem of locating and orienting a network of unattended sensors that has been deployed in a scene with unknown sensor locations and orientation angles, when no "anchor" nodes are present. Many localization problems assume that some nodes have known locations and propagate location information about other nodes using triangulation procedures. In our formulation, we do not require such anchor nodes, but instead assume prior probability density function for the nominal locations of a subset of the nodes. These nominal locations typically have high uncertainty, on the order of tens of meters. The self-calibration solution is obtained in two steps. Relative sensor locations are estimated using noisy time-of-arrival and direction-of-arrival measurements of calibration source signals in the scene, and absolute location calibration is obtained by incorporating prior nominal location knowledge. We consider a Bayes approach to the calibration problem and compute accuracy bounds on the calibration procedure. A maximum *a posteriori* estimation algorithm is shown to achieve the accuracy bound. Experiments using both synthetic data and field measurements are presented.

Keywords: sensor networks, localization, self-calibration, maximum *a posteriori estimation*, Cramér-Rao Bound

1. INTRODUCTION

Collaborative sensing has becoming increasingly important in many military and civilian applications.[1–4] The basic concept is to deploy a large number of low-cost, self-powered sensor nodes that acquire and process data. Typical sensors may include one or more acoustic microphones as well as seismic, magnetic, or imaging sensing capability. The collaborative sensing goals are to detect, track, and classify objects in the environment. An RF communication network links the sensors to one another and to a Central Information Processor, which in turn communicates with a higher-level information processing center (see Figure 1).

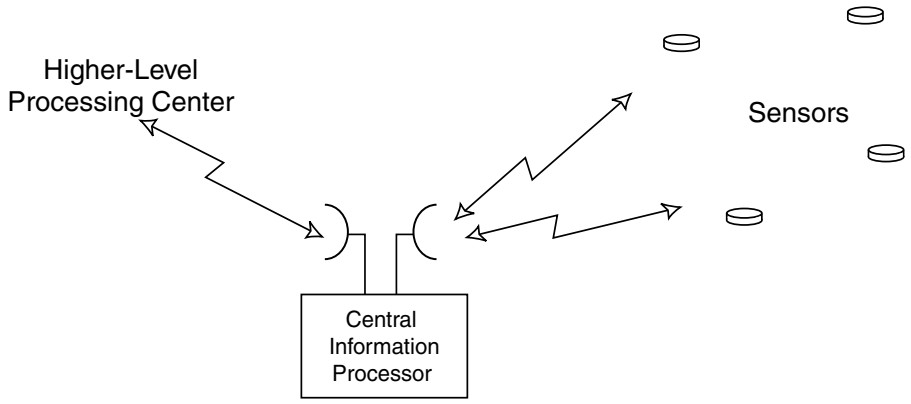

Figure 1. Sensor network architecture. A network of low-cost sensors are deployed in a region. Each sensor communicates to a local Central Information Processor, which relays information to a higher-level information processing center.

In order to fuse information acquired by individual sensors, one generally needs to know the location and orientation of each sensor. However, accurate prior knowledge of sensor locations and orientations is often not available. These sensors are placed in the field by persons, by an air drop, or by munition launch, and in each case the location and orientation of the placed sensor has high uncertainty. One could equip every sensor with a GPS and compass to obtain location and orientation information, but this adds to the expense and power requirements of the sensor and increases its susceptibility to detection or jamming. Thus, there is interest in developing methods to self-localize the sensor network with a minimum of additional hardware, processing, or communication.

Self-localization in sensor networks is an active area of current research.[5–7] Savvides *et al.*[8] consider iterative multilateration-based techniques, and Bulusu *et al.*[5,9] consider low-cost localization methods that employ beacon signals at known locations. Cevher and McClellan[10] consider sensor network self-calibration using a single acoustic source that travels along a straight line. Bearings-only localization methods[11] have also been considered.

Many self-calibration techniques require knowledge of the absolute locations of some nodes. In particular, the above techniques assume exact knowledge of a small number of "anchor nodes". Such information is often not available, or is costly to provide. In this paper we consider an approach to sensor network self-calibration using a combination of calibration sources and uncertain prior location information. The proposed self-calibration approach entails placing a number of signal sources in the same region as the sensors (see Figure 2). Neither the source locations nor their signal emission times are assumed to be known. Each source generates a signal that is detected by some of the sensors, and each sensor measures the time-of-arrival (TOA) and, if possible, direction-of-arrival (DOA) of the source signal. These measurements permit estimating the *relative* locations and orientations of the sensors.

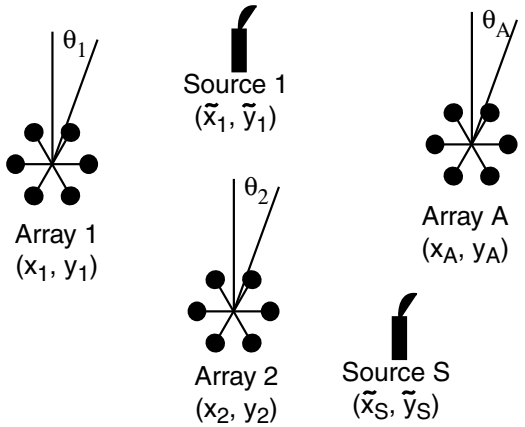

Figure 2. Sensor self-localization scenario. Sensors are placed in the field at unknown locations and orientations. Calibrations sources, also at unknown locations, generate signals time-of-arrival and direction-of-arrival measurements at the sensors are used for network self-calibration.

In addition, we assume uncertain prior location information for some of the sensors. Uncertain prior information about sensor location may be available in practice from aimpoints of munition-deployed sensors, or from aircraft location at the point of an air drop of a sensor. Although the aimpoints may be known, the actual sensor location typically has high uncertainty (on the order of tens of meters) from this aimpoint.

We develop maximum *a posteriori* (MAP) estimates of the absolute locations and orientations of the sensors from the prior location information along with the source-to-sensor TOA and DOA estimates. We also derive bounds on localization accuracy. We parameterize the localization problem as a sum of two parts: an absolute location and orientation of the entire sensor network, and the relative location and orientation of each sensor with respect to this absolute location. We compute the Cramér-Rao bound for the relative location parameters and derive the uncertainty of the MAP absolute location estimate. The relative location errors are typically (much) smaller than the prior location uncertainties.

Finally, we present experimental results on both synthetic data and field measurements to demonstrate the effectiveness of the proposed techniques.

2. SELF-LOCALIZATION PROBLEM FORMULATION

Assume we have a set of A sensors in a plane, each with unknown location $\{a_i = (x_i, y_i)\}_{i=1}^{A}$ and unknown orientation angle θ_i with respect to a reference direction (*e.g.*, North); see Figure 2. We consider the two-dimensional problem in which the sensors lie in a plane and the unknown reference direction is azimuth; an extension to the three-dimensional case is possible using similar techniques.

In the sensor field are also placed S point calibration source signals at locations $\{s_j = (\tilde{x}_j, \tilde{y}_j)\}_{j=1}^{S}$. The source locations are in general unknown. Each source emits a finite-length signal that begins at time t_j; the emission times are also in general unknown.

Define the parameter vector α containing the sensor and source unknown parameters:

$$\alpha = [\beta^T, \gamma^T]^T \quad (3(A+S)\times 1) \tag{1}$$

$$\beta = [x_1,\, y_1,\, \theta_1, \ldots, x_A,\, y_A,\, \theta_A]^T \quad (3A\times 1) \tag{2}$$

$$\gamma = [\tilde{x}_1,\, \tilde{y}_1,\, t_1, \ldots, \tilde{x}_S,\, \tilde{y}_S,\, t_S]^T \quad (3S\times 1) \tag{3}$$

Here, β contains the sensor unknown parameters and γ contains the unknown nuisance source signal parameters. Let $n_\alpha = 3(A+S)$ denote the length of α.

We assume that some prior information about α is available. Typically the prior information encodes the nominal locations and uncertainty of some of the sensors or sources. We express this this prior information probabilistically as a prior probability density function $f_0(\alpha)$. As an example, one can assume $f(\alpha)$ is Gaussian distributed with known mean and covariance, so

$$f_0(\alpha) = \mathcal{N}(\alpha_0, \Sigma_0) \tag{4}$$

where α_0 and Σ_0 are given.

2.1. Absolute and Relative Calibration

It is useful to express the location parameters in terms of a relative location and orientation with respect to the sensor-source centroid. To this end, define the centroid location and orientation, $\alpha_c(\alpha) = [x_c, y_c, \theta_c]^T$, by

$$x_c = \frac{1}{A+S}\left[\sum_{i=1}^{A} x_i + \sum_{j=1}^{S} \tilde{x}_j\right] \qquad y_c = \frac{1}{A+S}\left[\sum_{i=1}^{A} y_i + \sum_{j=1}^{S} \tilde{y}_j\right] \qquad \theta_c = \frac{1}{A}\sum_{i=1}^{A} \theta_i \tag{5}$$

We see that $\alpha_c = B\alpha$ for a $(3 \times n_\alpha)$ matrix B with elements defined in (5).

For any parameter vector α, define the corresponding relative location vector α_r by transforming the points in α such that $B\alpha_r = 0$. The transformation is defined as follows for a vector α with centroid $\alpha_c = [x_c, y_c, \theta_c]^T$. For the ith sensor with parameters (x_i, y_i, θ_i), which are the $3(i-1)+1$*st* to $3(i-1)+3$*rd* elements of α, the corresponding elements of α_r are given by

$$\begin{bmatrix} x_{i,r} \\ y_{i,r} \\ \theta_{i,r} \end{bmatrix} = \begin{bmatrix} \cos\theta_c & \sin\theta_c & 0 \\ -\sin\theta_c & \cos\theta_c & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_i - x_c \\ y_i - y_c \\ \theta_i - \theta_c \end{bmatrix} \tag{6}$$

Similarly, for the jth source with parameters $(\tilde{x}_j, \tilde{y}_j, t_j)$ (which are the $3(A+j-1)+1$*st* to $3(A+j-1)+3$*rd* elements of α, the corresponding elements of α_r are given by

$$\begin{bmatrix} \tilde{x}_{j,r} \\ \tilde{y}_{j,r} \\ t_{j,r} \end{bmatrix} = \begin{bmatrix} \cos\theta_c & \sin\theta_c & 0 \\ -\sin\theta_c & \cos\theta_c & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{x}_j - x_c \\ \tilde{y}_j - y_c \\ t_j \end{bmatrix} \tag{7}$$

From the above definitions, it is clear that there is a one-to-one relationship between a parameter vector α and the two vectors α_r and α_c. Note that t_j remains unchanged by the transformation. The inverse transformations are readily obtained from (6)–(7). We remark that even though the centroid α_c can be computed linearly ($\alpha_c = B\alpha$), the transformation of a vector α to obtain α_r is nonlinear because all locations have to be rotated by θ_c. We will use α and (α_r, α_c) interchangeably in the remainder of the paper.

2.2. Calibration Measurements

Each emitted source signal is detected by all of the sensors in the field, and each sensor measures the time-of-arrival and direction-of-arrival for that source. We denote the measured TOA at sensor i of source j as t_{ij} and the measured DOA as θ_{ij}. The DOA measurements are made with respect to a frame of reference local to the sensor. The times of arrival are measured with respect to a known, common time base. The time base can be established, for example, by using them RF communication network linking the sensors. The time base needs to be accurate to on the order of the sensor's time of arrival measurement uncertainty (which is 1 msec in the examples considered in Section 6).

The set of $2AS$ calibration measurements are gathered in a vector

$$X = \begin{bmatrix} \mathrm{vec}(T) \\ \mathrm{vec}(\Theta) \end{bmatrix}^T \quad (2AS \times 1) \tag{8}$$

where $\mathrm{vec}(M)$ stacks the elements of a matrix M columnwise and where

$$T = \begin{bmatrix} t_{11} & t_{12} & \dots & t_{1S} \\ t_{21} & t_{22} & \dots & t_{2S} \\ \vdots & \vdots & \ddots & \vdots \\ t_{A1} & t_{A2} & \dots & t_{AS} \end{bmatrix}, \qquad \Theta = \begin{bmatrix} \theta_{11} & \theta_{12} & \dots & \theta_{1S} \\ \theta_{21} & \theta_{22} & \dots & \theta_{2S} \\ \vdots & \vdots & \ddots & \vdots \\ \theta_{A1} & \theta_{A2} & \dots & \theta_{AS} \end{bmatrix} \tag{9}$$

Each sensor transmits its $2S$ noisy TOA and DOA measurements to a Central Information Processor (CIP), and these $2AS$ measurements, along with the measurement uncertainty and prior uncertainty $f_0(\alpha)$ form the information with which the CIP computes the sensor calibration. The communication cost to the CIP is low, and the calibration processing can be entirely performed by the CIP.

We denote the true TOA and DOA of source signal j at sensor i as $\tau_{ij}(\alpha)$ and $\phi_{ij}(\alpha)$, respectively, and include their dependence on the parameter vector α; they are given by:

$$\tau_{ij}(\alpha) = t_j + \|a_i - s_j\|/c \tag{10}$$

$$\phi_{ij}(\alpha) = \theta_i + \angle(a_i, s_j) \tag{11}$$

where $a_i = [x_i, y_i]^T$, $a_j = [\tilde{x}_j, \tilde{y}_j]^T$, $\|\cdot\|$ is the Euclidean norm, $\angle(\xi, \eta)$ is the angle between the points $\xi, \eta \in \mathcal{R}^2$, and c is the signal propagation velocity.

Each element of X has measurement uncertainty; we model the uncertainty as

$$X = \mu(\alpha) + E \tag{12}$$

where $\mu(\alpha)$ is the noiseless measurement vector whose elements are given by equations (10) and (11) for values of i, j that correspond to the vector stacking operation in (8), and where E is a random vector with known probability density function.

The self-calibration problem, then, is: given the measurement X, estimate β. The parameters in γ are in general unknown and are nuisance parameters that may also need to be estimated.

Significantly, the measurements in X provide information about the *relative* sensor and source locations but not about the centroid. Thus, X enables us to estimate α_r but provides no information about α_c.

3. BAYESIAN SELF-CALIBRATION

We consider α as a random vector with prior probability density function (pdf) $f_0(\alpha)$. The measurement vector informs us about α as quantified by the posterior probability $f(\alpha|X)$. From Bayes' rule,

$$f(\alpha|X) = \frac{f(X|\alpha) f_0(\alpha)}{f(X)} \tag{13}$$

We choose as our estimate maximum *a posteriori* (MAP) estimate, which is the value of α that maximizes the posterior probability density of α:

$$\hat{\alpha} = \arg\max_{\alpha} f(\alpha|X) = \arg\max_{\alpha} f(X|\alpha) f_0(\alpha) \tag{14}$$

Note that the measurements X are independent of absolute location; thus, $f(X|\alpha) = f(X|\alpha_r)$. Combining this fact with the negative logarithm of equation (14) yields

$$\hat{\alpha} = \arg\min_{(\alpha_r, \alpha_c)} [-\ln f(X|\alpha_r) - \ln f_0(\alpha)] \tag{15}$$

where $\alpha_c = [x_c, y_c, \theta_c]^T$ is a 3×1 centroid vector.

3.1. Gaussian Measurement Uncertainty

If we assume the measurement uncertainty E in equation (12) is Gaussian with zero mean and known covariance Σ_X, then

$$f(X;\alpha) = \frac{1}{(2\pi)^{AS}|\Sigma_X|^{1/2}} \exp\left\{-\frac{1}{2} Q(X;\alpha)\right\} \tag{16}$$

$$Q(X;\alpha) = [X - \mu(\alpha)]^T \Sigma_X^{-1} [X - \mu(\alpha)] \tag{17}$$

A special case is when the measurement errors are uncorrelated and the TOA and DOA measurement errors have variances σ_t^2 and σ_θ^2, respectively; equation (17) then becomes

$$Q(X;\alpha) = \sum_{i=1}^{A} \sum_{j=1}^{S} \left[\frac{(t_{ij} - \tau_{ij}(\alpha))^2}{\sigma_t^2} + \frac{(\theta_{ij} - \phi_{ij}(\alpha))^2}{\sigma_\theta^2} \right] \tag{18}$$

If we further assume that the prior information is Gaussian with mean α_0 and covariance Σ_0, then

$$f_0(\alpha) = \frac{1}{(2\pi)^{3(A+S)/2}|\Sigma_0|^{1/2}} \exp\left\{-\frac{1}{2}[\alpha - \alpha_0]^T \Sigma_0^{-1} [\alpha - \alpha_0]\right\} \tag{19}$$

Note that some elements of α may have no prior information, in which case the above pdf will be reduced in dimension accordingly.

For the case of Gaussian priors and Gaussian measurement errors, equation (15) yields

$$\hat{\alpha} = \arg\min_{(\alpha_r, \alpha_c)} [X - \mu(\alpha)]^T \Sigma_X^{-1} [X - \mu(\alpha)] + [\alpha - \alpha_0]^T \Sigma_0^{-1} [\alpha - \alpha_0] \tag{20}$$

Equations (15) and (20) both involve minimization of a nonlinear function of α; for the Gaussian case, (20) yields a nonlinear least squares minimization problem. In the next section we discuss various approaches to solve this minimization problem.

4. ALGORITHMS FOR SELF-CALIBRATION

First, we write (15) as

$$\hat{\alpha} = \arg\min_{(\alpha, c)} [J_1(\alpha_r) + J_2(\alpha_r, c)] \tag{21}$$

where $J_1(\alpha_r) = -\ln f(X|\alpha_r)$ and $J_2(\alpha_r, c) = -\ln f_0(\alpha)$.

In developing algorithms for minimizing (21), it is useful to consider separately the case for which some prior information on α is known with relatively high accuracy and the case for which the prior information has high uncertainty. For the former case, one can use the high-accuracy prior information to efficiently determine an initial α to form the starting point for the nonlinear minimization. For example, if we know the prior location of at least two sensors with high accuracy, then the techniques developed in Moses *et al.*[12] can be used to initialize the locations orientations, and emission times of the sources and sensors.

The more difficult case is the one for which the prior location information has high uncertainty. In this case, we form an initial estimate by making the following approximation: we assume that α_r that minimizes (21) is approximately equal to $\hat{\alpha}_r$, where

$$\hat{\alpha}_r = \arg\min_{(\alpha_r)} J_1(\alpha_r) \tag{22}$$

Next, consider the minimization of $J_1(\alpha_r)$. We proceed as follows. First, note that there is a one-to-one correspondence between a parameter vector α_r and a parameter vector α_1 for which $x_1 = y_1 = \theta_1 = 0$. This is not surprising, since in both cases there are three degrees of freedom removed from a parameter vector α. One can compute α_r from α_1 using equations similar to (6)–(7). Thus, one method for finding $\hat{\alpha}_r$ is to find

$$\hat{\alpha}_1 = \arg\min_{\alpha_1} J_1(\alpha_1) \tag{23}$$

where α_1 satisfies $x_1 = y_1 = \theta_1 = 0$, and then to convert $\hat{\alpha}_1$ to a corresponding $\hat{\alpha}_r$ using this one-to-one transformation. An algorithm for finding $\hat{\alpha}_1$ in (23) is given in Moses *et al.*[12]

Once $\hat{\alpha}_r$ is obtained from (22), finding $\hat{\alpha}$ that solves (21) remains. This generally involves a nonlinear minimization over the parameters in α. Initial estimates for α_r have already been computed, so we need to initialize the 3×1 constraint vector α_c. One way to initialize α_c is to find the translation components x_c and y_c to align the centroid of the sensor and source locations to centroid of the prior location information. For example, if the prior information $f_0(\alpha)$ is Gaussian with mean α_0, and covariance $\Sigma_0 = kI$ for some constant k, one finds that the ML estimate of the centroid given α_r is found by setting (x_c, y_c) to the centroid of the known sensor and source locations given in α_0. In practice, it may be that prior location information for only a subset of sensors or sources is known; in this case one would set the centroids equal only on the subset for which prior information is known. The remaining centroid parameter, θ_c, can be initialized by a coarse brute-force search on the interval $[0, 2\pi)$. For Gaussian prior locations with general Σ_0, the solution for x_c and y_c that minimizes (21) can be found in closed form as the solution to a 2×2 matrix equation.

Once an initial estimate of α_c is found, one can find $\hat{\alpha}$ from (21) using an iterative descent procedure. In our implementation we used the Matlab function `lsqnonlin`.

If the prior knowledge has high uncertainty, an approximate solution to (21), with lower computational cost, may be found as follows. If the prior knowledge has high uncertainty, then J_2 depends only weakly on α_r, so the optimal α_r vector is approximately given by the solution to (22). Then, we can solve for the remaining three parameters in α_c by inserting $\hat{\alpha}_r$ into equation (21) and finding $\hat{\alpha}_c$; that is:

$$\hat{\alpha}_c = \arg\min_{\alpha_c} \left[J_1(\hat{\alpha}_r) + J_2(\hat{\alpha}_r, c) \right] \tag{24}$$

As discussed above, if the prior uncertainty model is Gaussian, the translation parameters x_c and y_c can be found in closed form, and the minimization in (24) reduces to a search over the scalar parameter θ_c.

4.1. Uncertainty in Propagation Velocity

Thus far we have assumed that the signal propagation velocity, c, is known. In this section we discuss modifications to the above approach when there is some uncertainty in c.

First, we note that the TOA and DOA measurements provide no information for estimation of c. Uncertainty in c results in an uncertainty on the scaling of the overall source-sensor locations in α_r. Thus, if c is unknown, we append c to the set of unknowns in α_c.

For the case of Gaussian prior information with uncorrelated and equal variance uncertainty in the sensor and source x and y locations, equation (24) (where now $\alpha_c = [x_c, y_c, \theta_c, c]$ can be solved for in closed-form using a least-squares estimate.[13] Said another way, the conditional maximum likelihood estimate of c given α_r is obtained by solving a linear set of equations. It is interesting to note that a closed-form solution is possible when c is unknown but not when c is known.

5. ESTIMATION ACCURACY

In this section we establish the accuracy of both the maximum likelihood estimate of α_r for the case of no prior information, and for the MAP estimate of α that incorporates both the TOA/DOA measurements and the prior location information. We use the Fisher Information Matrix (FIM) and the Cramér-Rao Bound (CRB) as tools to analytically determine estimation accuracy. The Cramér-Rao Bound (CRB) gives a lower bound on the covariance of any unbiased estimate of α. It is a tight bound in the sense that $\hat{\alpha}_{ML}$ has parameter uncertainty given by the CRB for high measurement signal-to-noise ratio; that is, as $\max_i(\Sigma_X)_{ii} \to 0$. Thus, the CRB is a useful tool for analyzing calibration uncertainty. It is computed from the inverse of the Fisher Information Matrix when the inverse exists. In our case, the FIM for α_r is singular, and a pseudoinverse-based solution must be used. In the addition of prior information, we can compute the covariance of the MAP α estimate.

5.1. Estimation Accuracy of Relative Location

First, we establish a lower bound on the covariance of an unbiased estimator of α_r based on measurements X. In this derivation, we assume no prior knowledge on α; that is, we assume α is a deterministic unknown parameter vector. Since X provides no information about the centroid parameters, we find that the Fisher information matrix of α is singular; in fact, the null space of the Fisher information matrix generally has rank three, corresponding to the three degrees of freedom in an arbitrary translation and rotation of the sensor and source locations. However, the constraints imposed on α_r are such that its covariance estimate is nonsingular in all but degenerate cases. An example of a degenerate case is when all sensors are collinear and a calibration source is collinear with the sensors and not between any two sensors.

The CRB can be computed from the Fisher Information Matrix. The Fisher Information Matrix of α is given by[14]

$$I_\alpha = E\left\{ \left[\nabla_\alpha \ln f(T,\Theta;\alpha)\right] \left[\nabla_\alpha \ln f(T,\Theta;\alpha)\right]^T \right\}$$

For the Gaussian estimation error case, the partial derivatives are readily computed from equation (16); we find that

$$I_\alpha = [G'(\alpha_1)]^T \Sigma_X^{-1} [G'(\alpha_1)] \tag{25}$$

where $G'(\alpha_1)$ is the $2AS \times n_\alpha$ matrix whose ijth element is $\partial\mu_i(\alpha_1)/\partial(\alpha_1)_j$.

The Fisher Information Matrix is rank deficient due to the translational and rotational ambiguity in the self-calibration solution.[12] However, the CRB of α_r is finite, and given by[14, 15]:

$$C_{\alpha_r} = P[I_\alpha]^\dagger P \tag{26}$$

where $(\cdot)^\dagger$ denotes the Moore-Penrose pseudoinverse and where P is a projection matrix given by

$$P = I - M(M^T M)^{-1} M^T \tag{27}$$

$$M = \left[\begin{array}{cccc|cccc} M_1 & M_2 & \cdots & M_A & \tilde{M}_1 & \tilde{M}_2 & \cdots & \tilde{M}_S \end{array}\right]^T \tag{28}$$

$$M_i = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -y_i & x_i & 1 \end{bmatrix}, \qquad \tilde{M}_j = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -\tilde{y}_j & \tilde{x}_j & 0 \end{bmatrix} \tag{29}$$

5.2. Accuracy of Maximum A Posteriori Absolute Location Estimate

When prior information is available, the covariance of the MAP estimate $\hat{\alpha}$ from (21) is found as follows. Since the prior information and the measurement information are independent, the information matrix associated with both is the sum of each. This total information matrix is given by

$$I_\alpha^0 = \Sigma_0^{-1} + I_\alpha \tag{30}$$

where I_α is given by equation (25) and where we have assumed a Gaussian prior. The covariance of the MAP estimate is, for high SNR, given by the inverse of this matrix (if it exists), so

$$\text{cov}(\hat{\alpha}) = [I_\alpha^0]^{-1} \tag{31}$$

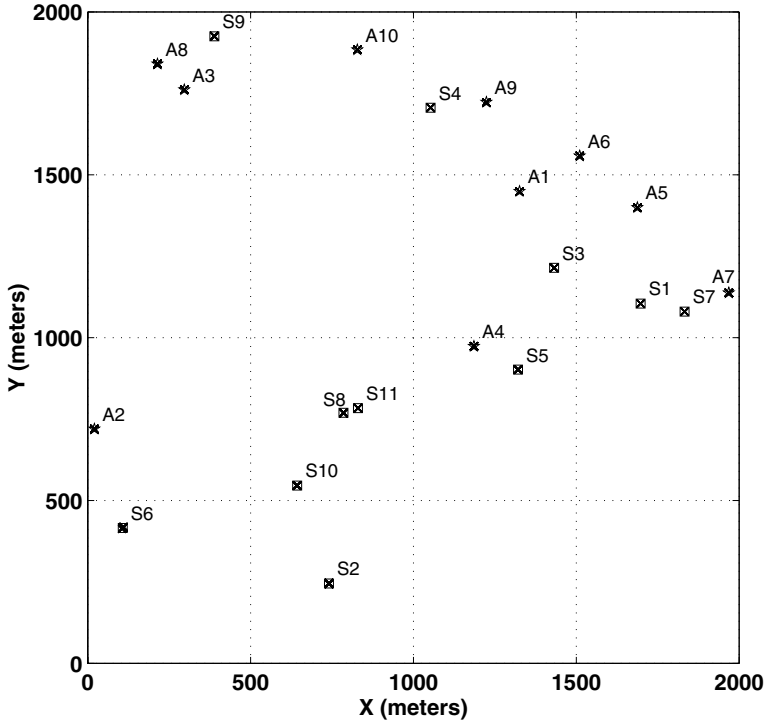

Figure 3. Top: Example scene showing ten sensors A1–A10 (stars) and eleven sources S1–S11 (squares). Also shown are the 2σ location uncertainty ellipses of the sensors and sources; these are on average less than 0.5 m in radius and show as small dots. Bottom: enlarged views near sensors A6 and A8, showing 2σ location uncertainty ellipses along with location estimates from 200 Monte-Carlo experiments.

6. EXPERIMENTAL RESULTS

This section presents numerical examples of the self-calibration procedure. First, we present a synthetically-generated example consisting of ten sensors and eleven sources placed randomly in a 2 km$\times$2 km region. Second, we present results from field measurements using acoustic sources and sensors.

6.1. Synthetic Data Example

We consider a case in which ten sensors are randomly placed in a 2 km$\times$2 km region. In addition, eleven sources are randomly placed in the same region. In both cases the nominal locations of the sensors and sources are first randomly chosen. Then, the actual locations are chosen as random Gaussian perturbations from the nominal locations, where the perturbation is found by adding Gaussian noise with zero mean and standard deviation of 10 meters to each location. The sensor orientations and source emission times are also randomly chosen. Figure 3 shows the actual locations of the sensors and sources.

Next, TOA and DOA measurements are simulated. We assume every sensor detects each source emission and measures the TOA and DOA of the source. The measurement uncertainties are Gaussian with standard deviations of $\sigma_t = 1$ msec for the TOAs and $\sigma_\theta = 3^\circ$ for the DOAs. Neither the locations nor emission times of the sources are assumed to be known.

Figure 3 also shows the two standard deviation (2σ) location uncertainty ellipses for the relative location estimates of both the sources and sensors. The ellipses are obtained from the 2×2 covariance submatrices of the CRB in equation (26) that correspond to the relative location parameters in α_r of each sensor or source. These ellipses appear as small dots in the figure; enlarged views for two sensors are also shown.

The results of the maximum likelihood estimation of relative location are also shown in Figure 4. The 'x' marks show the ML location estimates from 200 Monte-Carlo experiments with randomly-generated DOA and TOA measurements. The

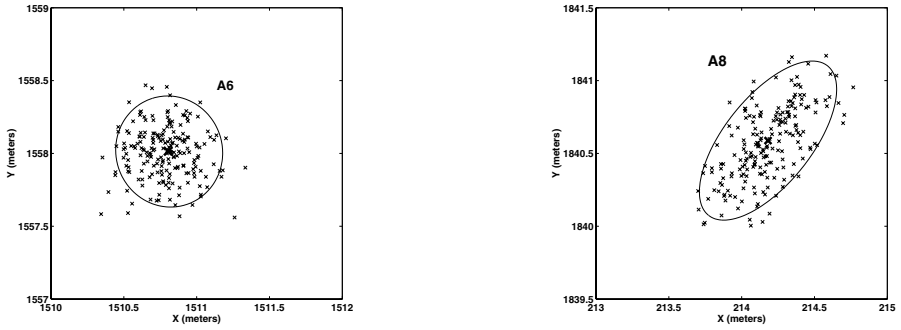

Figure 4. Two standard deviation location uncertainty ellipses for sensors $A6$ and $A8$ from Figure 3.

Table 1. Average 2σ uncertainty radius for relative and absolute sensor calibration for the example presented

	Theoretical Value	Monte-Carlo Simulation
Relative Location	0.47 m	0.43 m
Absolute Location	5.19 m	5.09 m

DOA and TOA measurement errors were drawn from Gaussian distributions with zero mean and variances of $\sigma_t = 1$ msec and $\sigma_\theta = 3^\circ$, respectively. The ellipse shows the 2-standard deviation (2σ) relative location uncertainty region as predicted from the CRB. We find good agreement between the CRB uncertainty predictions and the Monte-Carlo experiments, which demonstrates the statistical efficiency of the ML estimator for this level of measurement uncertainty.

We summarize the ten location uncertainty ellipses into a single quantity by first defining the sensor or source 2σ *uncertainty radius* as the radius of a circle whose area is the same as the area of the 2σ location uncertainty ellipse. The 2σ uncertainty radius for each sensor or source is computed as the geometric mean of the major and minor axis lengths of the 2σ uncertainty ellipse. We then average these uncertainty radii to obtain a single average location error quantity. Table 1 shows the average 2σ uncertainty ellipse radius for the ten sensors, as computed from the CRB, and also the estimated average uncertainty ellipse computed from Monte-Carlo experiments. The Monte-Carlo results are found by estimating α_r for each of the 200 Monte-Carlo experiments. We also compute the true α_r vector by transforming the true α vector so that its centroid is zero. In each Monte-Carlo experiment we compute the average distance from each sensor to its corresponding true relative location, then average this distance over all Monte-Carlo experiments. We see that the results agree, and the average uncertainty radius is approximately 0.45 m for this example. Not shown in the figure are the sensor orientation angles and their estimates; the 2σ angle error bound, predicted from the CRB, is 1.8° for each sensor, and the Monte-Carlo estimates are close to this value as well.

We also estimate the absolute location and the average location error. In this case we compute $\hat{\alpha}$ for each of the 200 Monte-Carlo experiments and then compute the average distance from the true locations as above. We compare this average 2σ error distance to the corresponding average error computed from $\mathrm{cov}(\hat{\alpha})$ in equation (31). These results are also shown in Table 1. We see that the absolute 2σ error is about 5 meters, approximately 1/4 the prior 2σ uncertainty of 20 meters for each sensor. The theoretical variance of 5.19 m in the table is typical; in another simulation we computed the theoretical 2σ absolute distance error for 200 choices of true sensor locations and found the average of these 200 distances to be 5.17 m.

This experiment shows that relative location calibration of the sensors to errors well less than 1 meter is possible. In addition, even with weak prior information of sensor locations, as might be available from aimpoints for sensor placement, absolute location uncertainties on the order of 5 meters can be obtained. It is significant that the relative location errors are small — this means that beamforming or source location and tracking can be expected to have high accuracy. The absolute location corresponds to uncertainty only in the translation or rotation of the entire sensor network. While this leads to a corresponding uncertainty in, for example, a target being tracked by the network, the uncertainty lies largely in

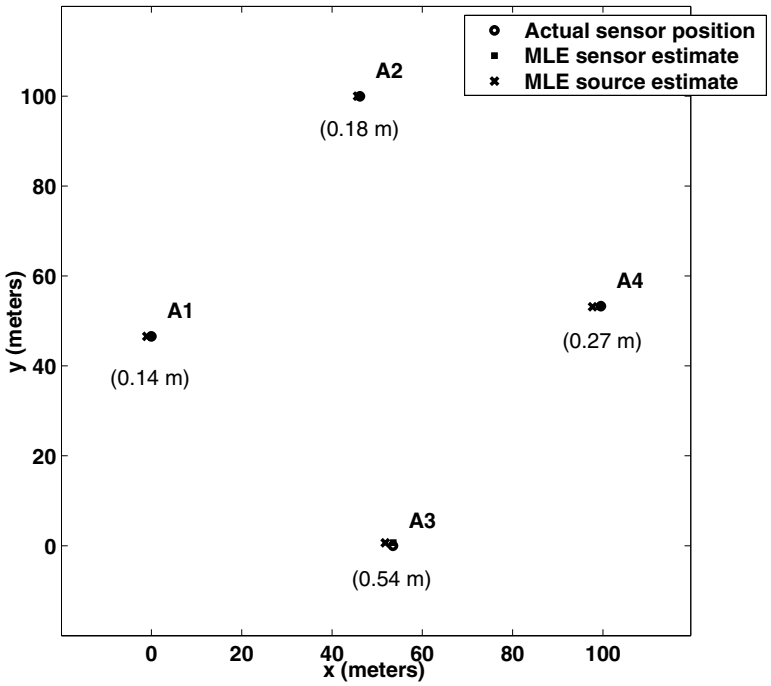

Figure 5. Actual and estimated sensor locations, and estimated source locations, using field test data. Error distances between actual and estimated sensor locations are shown in parentheses.

the translation and rotation of the track. Other available information may be used to further correct for this location and translation error. If a track is found to be slightly translated and rotated from a known road, for example, the absolute orientation sensor network can be correspondingly adjusted.

6.2. Field Test Results

We present the results of applying the auto-calibration procedure to an acoustic source calibration data collection conducted during the DUNES test at Spesutie Island, Aberdeen Proving Ground, Maryland in September 1999. In this test, four acoustic sensors are placed at known locations 60-100 m apart as shown in Figure 5. Four acoustic source signals are also used. Exact ground truth locations of the calibration sources are not known, but it is known that each source is within approximately 1 m of a sensor. Each source signal is a series of bursts in the 40-160 Hz frequency band. Time-aligned samples of the sensor microphone signals are acquired at a sampling rate of 1005.53 Hz. Times of arrival are estimated by cross-correlating the measured microphone signals with the known source waveform, and finding the peak of the correlation function. Only a single microphone signal is available at each sensor, so while TOA measurements are obtained, no DOA measurements are available. Figure 5 shows the ML estimates of the relative sensor and source location, as compared to the known relative locations of the sensors. The location errors of sensors A1, A2, A3, and A4, are 0.14 m, 0.18 m, 0.54 m and 0.27 m, respectively, for an average error of 0.28 m. In addition, the source location estimates are within 1 m of the sensor locations, consistent with our ground truth information.

7. CONCLUSIONS

We have presented a procedure for calibrating the locations and orientations of a network of sensors. The calibration procedure uses source signals that are placed in the scene and computes sensor and source unknowns from estimated time-of-arrival and direction-of-arrival estimates obtained for each source-sensor pair. These measurements can be used to compute a maximum likelihood estimate of the relative sensor and calibration source locations, along with the relative orientation of the sensors. Absolute sensor calibration is obtained by using these measurements along with probabilistic prior location information of some of the sources. A maximum *a posteriori* calibration solution is proposed in this case.

An analytical expression for the Cramér-Rao lower bound on the relative sensor calibration error covariance matrix is presented. In addition, a high measurement SNR expression for the posterior error covariance of the absolute calibration method is derived. Experimental results on synthetic and measured data validate the theoretical expressions and show that two-standard-deviation relative location accuracies of less than 1 meter can be obtained. Absolute location calibration two-standard-deviation uncertainty of approximately 5 meters was obtained for the cases considered.

The calibration procedure requires low sensor communication and has reasonable computational cost. The algorithms require low communication overhead as each sensor needs to communicate only two real values to the CIP for each source signal it detects. Computation of the calibration solution takes place at the CIP. For the synthetic examples presented the calibration computation takes on the order of 10 seconds using Matlab on a standard personal computer. For the field test data, computation time was less than 1 second.

ACKNOWLEDGEMENT

This material is based in part upon work supported by the U.S. Army Research Office under Grant No. DAAH-96-C-0086 and Batelle Memorial Institute under Task Control No. 01092, and in part through collaborative participation in the Advanced Sensors Consortium sponsored by the U.S. Army Research Laboratory under the Federated Laboratory Program, Cooperative Agreement DAAL01-96-2-0001. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the U.S. Army Research Office, the Army Research Laboratory or the U.S. Government.

REFERENCES

1. D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the world with wireless sensor networks," in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, **4**, pp. 2033–2036, (Salt Lake City, UT), May 7–11 2001.
2. G. Pottie and W. Kaiser, "Wireless integrated network sensors," *Communications of the ACM* **43**, pp. 51–58, May 2000.
3. N. Srour, "Unattended ground sensors a prospective for operational needs and requirements," tech. rep., Army Research Laboratory, October 1999.
4. "Collaborative signal and information processing in microsensor networks," *IEEE Signal Processing Magazine* **19**, March 2002.
5. N. Bulusu, J. Heidemann, and D. Estrin, "GPS-less low-cost outdoor localization for very small devices," *IEEE Personal Communication* **7**, pp. 28–34, October 2000.
6. C. Savarese, J. Rabaey, and J. Beutel, "Locationing in distributed ad-hoc wireless sensor networks," in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, **4**, pp. 2037–2040, (Salt Lake City, UT), May 7–11 2001.
7. J. C. Chen, K. Yao, and R. E. Hudson, "Source localization and beamforming," *IEEE Signal Processing Magazine* **19**, pp. 30–39, March 2002.
8. A. Savvides, C. C. Han, and M. B. Srivastava, "Dynamic fine-grained localization in ad-hoc wireless sensor networks," in *Proceedings of the International Conference on Mobile Computing and Networking (MobiCom) 2001*, (Rome, Italy), July 2001.
9. N. Bulusu, D. Estrin, L. Girod, and J. Heidemann, "Scalable coordination for wireless sensor networks: Self-configuring localization systems," in *Proceedings of the Sixth International Symposium on Communication Theory and Applications (ISCTA '01)*, July 15–20 2001.
10. V. Cevher and J. H. McClellan, "Sensor array calibration via tracking with the extended kalman filter," in *Proceedings of the Fifth Annual Federated Laboratory Symposium on Advanced Sensors*, pp. 51–56, (College Park, MD), March 20–22 2001.
11. L. Kaplan, L. Q. Le, and P. Molnár, "Maximum likelihood methods for bearings-only target localization," in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, **5**, pp. 3001–3004, (Salt Lake City, UT), May 7–11 2001.
12. R. L. Moses, D. Krishnamurthy, and R. Patterson, "A self-localization method for wireless sensor networks," *Eurasip Journal on Applied Signal Processing, Special Issue on Sensor Networks* . (submitted November 2001).

13. P. B. van Wamelen, Z. Li, and S. S. Iyengar, "A fast expected time algorithm for the point pattern matching problem," *Pattern Recognition* , November 2000. (submitted).
14. H. L. Van Trees, *Detection, Estimation, and Modulation Theory: Part I*, Wiley, New York, 1968.
15. P. Stoica and T. L. Marzetta, "Parameter estimation problems with singular information matrices," *IEEE Transactions on Signal Processing* **49**, pp. 87–90, January 2001.