A FULLY DISTRIBUTED METHOD FOR ACOUSTIC LOCALIZATION WITH SENSOR NETWORKS

By

Stephen Michael Williams

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

Mechanical Engineering

August, 2005

Nashville, Tennessee

Approved:

Kenneth Frampton

Nilanjan Sarkar

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation CAREER Program under Grant No. 0134224.

I would like to thank members of my Thesis Committee, Professor Nilanjan Sarkar and Professor Eric Barth for their time spent reviewing this paper. For their contribution to this work I would also like to thank Ph.D. candidates Isaac Amundson and Peter Schmidt. Finally, I would like to thank Professor Ken Frampton, the chairman of my Thesis Committee, for his professional guidance and support throughout my career at Vanderbilt.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER I

INTRODUCTION

The past two decades have seen a great deal of research in the area of distributed sensor networks. Aided by the availability of small, inexpensive microsensors, researchers have been investigating the usefulness of distributed sensor networks for a variety of applications. These applications include military surveillance, habitat monitoring, distributed robotics, air traffic control, and building security.

Distributed systems consisting of numerous "nodes", each possessing a microprocessor, power supply, wired or wireless communication capability, sensors and signal conditioning circuitry, face a number of technical challenges. Among these are energy and bandwidth constraints, ad hoc networking, collaborative information processing, message routing, and security. Performing collaborative information processing over a network is related to distributed data fusion and introduces additional technical challenges to the development of distributed sensor networks. The degree of information sharing among network nodes and the manner in which nodes fuse information from other nodes effects performance of the sensor network. Including information from more sensors typically results in better system performance, however, this also increases communication demands on the system. Communication requires energy, so high communication volume reduces the lifetime of sensor networks with finite energy resources. Therefore, researchers must consider the tradeoff between system performance and energy demands when designing a distributed sensor network [1, 2].

Decentralized algorithms, also called localized algorithms, may be used in place of centralized algorithms to reduce energy demands on the system that are introduced by high communication requirements. These decentralized algorithms are performed by nodes throughout the network, and incorporate information from nodes near those performing the computations. When nodes are added to a centralized system, communication requirements increase dramatically as all information must be relayed to a central controller. A system utilizing a decentralized algorithm will not see this dramatic increase in communications cost as nodes are introduced, and is therefore more scalable than its centralized counterpart. Additionally, localized algorithms are attractive because they are robust to network changes and node failures [1]-[3].

Acoustic Source Localization

Identifying and locating targets of interest are important tasks in military surveillance, habitat monitoring, and building security. Some targets (e.g. a vehicle, sniper, or other intruder) may possess an acoustic signature that can be used to identify the target. Once a target is identified, observation of acoustic events generated by the target may be used to locate its position. The process of determining the location of an acoustic source relative to some reference frame is known as acoustic source localization.

A distributed sensor network consisting of hundreds or thousands of nodes can be used to passively monitor the acoustical phenomena occurring within a large area. Sources in the near-field (geographical area encompassed by the network) can be localized with knowledge of the time-differences-of-arrival (TDOA) of the acoustic source at multiple nodes, and the speed of sound in the medium in which the acoustic source is present. Localizing sources in the far-field must be accomplished by methods other than those based on TDOA, such as, beamforming, direction-of-arrival (DOA), and energy based methods [4]-[9].

Communication costs in the large-scale sensor network described above will be astronomical if localization algorithms are implemented in a centralized fashion. It is therefore desirable to split the network into smaller groups consisting of a few nodes spaced closely together. Each group will have a leader who collects data from group members and then uses the information to estimate source positions. Performing source localization in this decentralized manner will lengthen the life-expectancy of the sensor network, and remove the need for a powerful central controller. Implementing a fully distributed algorithm will also mean that additional nodes may be seamlessly integrated into the system, and that the system will be resilient to node failures.

Others have developed distributed sensor networks for the purpose of acoustic localization. Lendeczi et al. developed a distributed system for battlefield surveillance using 56 Mica2 Motes manufactured by Crossbow Technology Inc. [8]. The system performs a grid search to find the location from which a rifle was fired using time-of-arrival (TOA) information of the muzzle blast and shock wave generated by the projectile fired from the rifle. TOA information is relayed to a base station where the location of the source is determined.

Chen et al. present a distributed sensor network capable of locating continuous near and far-field sources using approximated maximum likelihood and TOA methods in

[5]. This system relays information of the acoustic signal to a laptop where the source location is estimated.

Brooks, Ramanathan, and Sayeed discuss the challenges of implementing a distributed sensor network capable of tracking an acoustic source through a sensor field using a decentralized localization algorithm. They present solutions to provide efficient network techniques and collaborative signal processing in [7].

Acoustic Self-Localization

Many applications of distributed systems require nodes in the network to have knowledge of both their geographical position, and the positions of other nodes in the network. The task of determining node locations is an important aspect of distributed systems which has not previously been discussed, and is referred to as self-localization. In planned networks, the topology of the distributed network is usually known *a priori*, and this information is passed to nodes prior to system use [2]. For some applications, such as battlefield surveillance, nodes may be distributed randomly and *a priori* information of node locations is not available. Self-localization by GPS may be too expensive or not feasible (e.g. indoors) for a given network and other methods of self-localization must be employed.

"The most accurate methods for self-localization are based on either acoustic or RF wave propagation and time-of-flight measurements. While RF based localization is promising, there are several reasons why acoustic techniques may be more desirable. First of all the relatively low acoustic wave speed results in less expensive hardware as compared to high-bandwidth RF hardware. The lower wave speed also means that the system will be less sensitive to errors in time-of-flight measurements, node clock synchronization and other timing errors. However acoustic techniques also have two distinct disadvantages over RF techniques. First, acoustic waves are sensitive to environmental factors such as temperature, humidity and pressure resulting in non-uniform wave speeds over the area of interest. Second, acoustic waves suffer far greater attenuations through walls and other physical objects rendering it nearly ineffective when such obstacles exist." [10]

Researchers have successfully used acoustic wave propagation to perform self-localization in distributed sensor networks. Chen et al. [5], and Girod and Estrin [11] perform self-localization from time-of-flight measurements between nodes. Network nodes produce an acoustic chirp and simultaneously send a network message containing the time at which the chirp was produced to other nodes in the network. Nodes detect the TOA of the chirp with a matched filter and use this information and the time at which the chirp was produced to calculate the time-of-flight of the acoustic wave. The distance between nodes can then be found directly with knowledge of the speed of sound. By progressing through all nodes the range between each node and every other node were calculated and an optimization algorithm was used to calculate absolute coordinates. The system produced by Girod and Estrin was successful in self-localizing nodes to within about 11cm [11]. In this approach, although the source locations where not known, they were constrained to be the same as the node locations.

"Although the approach taken by Girod and Estrin, in which each node emits an acoustic signal to determine ranging, is very accurate there will be applications in which this approach would be undesirable. An example would be the case where the user does not want the sensor node locations to be revealed to observers in the area. If the nodes themselves are chirping then the security of the nodes would be compromised. Furthermore, acoustic emission is a power-hungry process if the signal must propagate very large distances. If the nodes themselves must be responsible for delivering this power their life expectancy will be significantly limited." [10]

Dosso et al. found the position of hydrophone array elements by solving linearized TOA equations with an iterative regularized inversion technique [12]. Acoustic sources (imploding light bulbs) were not attached to array elements for Dosso's work, and the time at which the source was produced was unknown. Instead, the locations of the sources were known. This method can be adapted to land-based sensor networks such as that consisting of hundreds of nodes distributed around an area of interest. A GPS equipped acoustic source could be moved throughout the sensor network, simultaneously emitting a sound and communicating its position to nodes in the network. Nodes would use the arrival time of the acoustic source and the location from which it was emitted to perform self-localization.

System benefits, such as lowered power consumption and scalability, achieved by implementing source localization algorithms in a decentralized manner may also be realized when performing self-localization.

Objective of This Work

The objective of this work is to produce a distributed sensor network from commercial off-the-shelf (COTS) products that is capable of self-localization and source localization via decentralized localization algorithms. Chapter II is a paper submitted to be published by *The Journal of Applied Acoustics* which reviews the importance of source localization; presents the theory behind the TDOA method for source localization; introduces an experimental distributed platform created for this work; and presents experimental analysis of the methods of source localization considered. Chapter III discusses the theory behind the TDOA method of self-localization, and presents experimental analysis of the methods considered.

CHAPTER II

SOURCE LOCALIZATION

The availability of miniature, low-power sensing devices has inspired the development of distributed sensor networks for a variety of applications [1, 2]. Among these are acoustic source localization and tracking, two integral tasks in surveillance, and habitat monitoring. Of interest here are systems that can effectively monitor a large area. This is a task naturally suited to a distributed sensor network made up of numerous "nodes" each possessing a microprocessor, power supply, wired or wireless communication capability, sensors and signal conditioning circuitry.

Distributed sensor networks face a number of technical challenges such as energy and bandwidth constraints, ad hoc networking, collaborative information processing, message routing, and security. The degree of information sharing among network nodes and the manner in which nodes fuse information from other nodes effects performance of the sensor network. Including information from more sensors typically results in better system performance, however; this also increases communication demands on the system, effectively reducing the lifetime of sensor networks with finite energy resources. Therefore, researchers must consider the tradeoff between system performance and energy demands when designing a distributed sensor network

The scenario that motivates the current work involves numerous (10s, 100s, or 1000s) of inexpensive nodes distributed around an area of interest with the objective of locating acoustical phenomena in that area. These nodes may be placed by hand;

dropped by UAV or UGV; or distributed by some other means, and the positions of nodes in the network are determined prior to system use. To effectively monitor the area, the sensor network must sense and identify acoustic sources, determine their position, and then relay information to appropriate locations.

Traditional sensor networks employ a centralized controller which directs all network activity and performs any computation. Using a centralized localization algorithm for the large-scale sensor network described above would place significant communication demands on the system as all source information would have to be relayed to the central controller. Decentralized algorithms, also called localized algorithms, may be used in place of centralized algorithms to reduce energy demands on the system which are introduced by high communication requirements. These decentralized algorithms are performed by nodes throughout the network, and incorporate information from nodes near those performing the computations. When nodes are added to a centralized system communication requirements increase dramatically as all information must be sent to a central controller. A system utilizing a decentralized algorithm will not see this dramatic increase in communication cost as nodes are introduced, and is therefore more scalable than its centralized counterpart. Additionally, localized algorithms are attractive because they are robust to network changes and node failures [1]-[3].

Source localization using sensor arrays has been studied for many years and currently sees application in radar, sonar, and wireless communication [4]. Recently, a number of distributed source localization systems have been proposed to estimate the position and/or the direction of arrival (DOA) of an acoustic source from phase and time-of-arrival (TOA) information measured by multiple receivers [5]-[9]. These systems vary in their method of localization (beamforming, DOA, TOA), as well as the nature of the sensor network that they employ (centralized, distributed), and there is generally a tradeoff between accuracy and computational demands. Lendeczi et al. present a centralized system for battlefield surveillance consisting of 56 Mica2 Motes manufactured by Crossbow Technology Inc. [8]. The system performs a grid search to find the location from which a rifle was fired using TOA information of the muzzle blast and shock wave generated by the projectile fired from the rifle. Chen et al. present a distributed sensor network capable of locating continuous near and far-field sources using approximated maximum likelihood and TOA methods in [5]. These systems relay information of the acoustic signal to a base station or laptop where the source location is estimated.

Estimating source position from a set of measured TOA data represents a nonlinear inverse problem in which the accuracy of the solution is dependent upon the solution method. The general least squares (LS) solution to the inverse problem seeks to minimize misfit to the measured data but may result in significant error when the problem is ill-posed. Dosso et al. present an iterative solution technique to localize elements of a horizontal line array which avoids the ill-posed nature of the problem [12]. The method is formulated to achieve misfit to the measured data consistent with estimated uncertainties of the data, while including independent information about the solution, known as *a priori* information, to the inversion. Node positions are treated as known parameters when solving the inversion process using LS, however; error in node position measurements may introduce significant error to the solution. This error may be reduced by treating both node and source positions as unknown parameters in the inversion, and then including node position estimates as *a priori* information in the iterative solution process.

A distributed sensor network utilizing a decentralized localization algorithm will maintain advantages in scalability, fault tolerance, and communication requirements over a comparable centralized system. Accordingly, the objective of this work is to produce a scalable distributed sensor network from commercial off-the-shelf (COTS) products that can localize a single source. Solutions of TDOA equations by LS and regularized inversion (RI) are presented in the Section II of this paper. A distributed system consisting of PC/104 modules is then introduced in Section III. The system is used to localize sources in the near and far-fields by both LS and RI methods, and results of these experiments are presented in Section IV. Conclusions based on these results are provided in Section V.

TDOA Formulation

Estimating the source location from measured TOAs represents an ill-posed, nonunique, nonlinear inverse problem. This paper will consider two methods for solving this inverse problem, and discuss the effectiveness of each method.

The first solution technique that will be considered is that formulated by Mahajan and Walworth in [13]. The formulation is based on differences in time-of-flight from a single source to multiple sensors. If we define the TOA of the acoustic source at the i^{th} node as T_i, the time-difference-of-arrival (TDOA) between a reference node and any other node is

$$T_{1i} = T_i - T_1 \quad (1)$$

where T_1 is the absolute arrival time of the acoustic source at the reference node. We can write the two-dimensional set of nonlinear equations representing the distances between the source and *N* nodes as

$$\begin{aligned} d^2 &= (x_1 - u)^2 + (y_1 - v)^2 \\ (d + cT_{12})^2 &= (x_2 - u)^2 + (y_2 - v)^2 \\ &\vdots \\ (d + cT_{1N})^2 &= (x_N - u)^2 + (y_N - v)^2 \end{aligned} \quad (2)$$

where *d* represents the distance from node 1 to the source (u,v), (x_i, y_i) is the position of the *ith* node, and *c* is the speed of sound. Expanding these equations and substituting the first equation for d^2 into the remaining equations results in the following linearized system of *N*-1 equations

$$\begin{bmatrix} 2x_1 - 2x_2 & 2y_1 - 2y_2 & -2cT_{12} \\ 2x_1 - 2x_3 & 2y_1 - 2y_3 & -2cT_{13} \\ & \vdots & \\ 2x_1 - 2x_N & 2y_1 - 2y_N & -2cT_{1N} \end{bmatrix} * \begin{bmatrix} u \\ v \\ d \end{bmatrix} = \begin{bmatrix} c^2T_{12}^2 + x_1^2 + y_1^2 - x_2^2 - y_2^2 \\ c^2T_{13}^2 + x_1^2 + y_1^2 - x_3^2 - y_3^2 \\ \vdots \\ c^2T_{1N}^2 + x_1^2 + y_1^2 - x_N^2 - y_N^2 \end{bmatrix} \quad (3)$$

This formulation (speed of sound is considered constant) can be solved when four nodes hear the acoustic source, and the system is overdetermined when more than four nodes are present. The above analysis can easily be extended to the three-dimensional case, requiring a fifth receiver to hear the source [13]. The estimate of the source location from (3) is then found by the method of least squares (LS). The system matrix in (3) may be ill-conditioned when relative time delays are approximately equal, resulting in significant error in the source position estimate. This is more likely to occur when the system is exactly determined, thus it is desirable to include more than four nodes in the problem formulation [6].

The second method used in this work to estimate source position was adapted from a technique formulated by Dosso et al., to localize horizontal line array elements [12]. Using the regularization method, an iterative linearized inversion is developed below which results in a stable solution to the ill-posed inverse problem.

Writing two-dimensional equations for the TOA of a single source at each node results in the following:

$$\begin{aligned} T_1 &= \frac{\sqrt{(x_1 - u)^2 + (y_1 - v)^2}}{c} + \gamma \\ &\vdots \\ T_N &= \frac{\sqrt{(x_N - u)^2 + (y_N - v)^2}}{c} + \gamma \end{aligned} \qquad (4)$$

where γ is the source instant. Note that this formulation may be extended to three-dimensions by including a third spatial coordinate in (4). Node locations were considered to be known in the previous formulation but will instead be treated as unknown, but well estimated, here. The system in (4) is a set of N nonlinear equations in M unknowns where M is equal to $2N + 3$ (two spatial unknowns per node in addition to the source location and source instant). This set of equations written in general vector form is

$$\boldsymbol{T = F(m)} \qquad (5)$$

where $\boldsymbol{T}$ represents the vector of arrival times at each node, $\boldsymbol{F}$ is the forward mapping matrix of the equations presented in (4), and $\boldsymbol{m}$ is the vector of model parameters [$x_1, y_1, \ldots, x_N, y_N, u, v, \gamma$]. A local linearization of the system is obtained by performing the Taylor series expansion of $\boldsymbol{T = F(m_o + \delta m)}$ about an arbitrary starting model $\boldsymbol{m_o}$. Disregarding higher order terms, the linearized system can be written as

$$\boldsymbol{T = F(m_o) + \mathbf{J}\delta m} \qquad (6)$$

where $\boldsymbol{\delta m}$ is the model perturbation, and **J** is the Jacobian matrix consisting of elements $J_{kl} = \partial F_k(\boldsymbol{m_o})/\partial m_l$. Defining the residual $\boldsymbol{\delta T = T - F(m_o)}$, (6) can be written as

$$\mathbf{J}\boldsymbol{\delta m} = \boldsymbol{\delta T} \tag{7}$$

which is a set of linear equations that can be solved for the model perturbation $\boldsymbol{\delta m}$. The model perturbation may be written as $\boldsymbol{\delta m = m - m_o}$ so that (7) can be rewritten in terms of the actual model, $\boldsymbol{m}$, as

$$\mathbf{J}\boldsymbol{m} = \boldsymbol{\delta T} + \mathbf{J}\boldsymbol{m_o} \tag{8}$$

This linearized inverse problem can be solved for $\boldsymbol{m}$, however; $\boldsymbol{m}$ may not adequately reproduce the measured data because nonlinear terms were neglected. If the model $\boldsymbol{m}$ does not reproduce the measured data, the starting model $\boldsymbol{m_o}$ is updated ($\boldsymbol{m_o} \leftarrow \boldsymbol{m}$), and the process repeated iteratively until an acceptable solution is found or the iterations converge.

The LS solution of the overdetermined system of linear equations given in (8) is found by minimizing the χ^2 misfit defined by

$$\chi^2 = |\boldsymbol{G(Jm-(\delta T} + \mathbf{J}\boldsymbol{m_o}))|^2 \tag{9}$$

where $\boldsymbol{G}$ is a diagonal weighting matrix, $\boldsymbol{G} = \mathrm{diag}[1/\sigma_1, \ldots, 1/\sigma_N]$. Note that (9) is the misfit of the model to the linearized inverse problem, and σ_i is the standard deviation of the TOA measurements at each node assuming that the error in TOA measurement can be represented as an independent Guassian-distributed random variable with mean of zero. The objective of this regularization is to formulate a unique, stable inversion by specifically including *a priori* information about the solution. This may be accomplished by minimizing an objective function Φ which combines the χ^2 term representing the data misfit, and a regularizing term that imposes the condition that the solution $\boldsymbol{m}$ resemble a prior estimate $\boldsymbol{m'}$ that includes *a priori* information.

$$\Phi = |\boldsymbol{G(Jm-(\delta T + \mathbf{J}m_o))}|^2 + \mu|\boldsymbol{H(m-m')}|^2 \tag{10}$$

The weighting matrix $\boldsymbol{H}$ in (10) is referred to as the regularization matrix. The parameter μ is a Lagrange multiplier which controls the relative importance of the data misfit and the *a priori* information. Estimates of the node positions were available during experimentation and were included in $\boldsymbol{m'}$. The regularization matrix was then held to be $\boldsymbol{H} = \mathrm{diag}[1/\sigma_{x1}, 1/\sigma_{y1}, \ldots, 1/\sigma_{xN}, 1/\sigma_{yN}, 0, 0, 0]$, where σ_{xi} and σ_{yi} are the standard deviations of node coordinates (again assuming that the error in node position could be represented as an independent, Gaussian-distributed random variable with a mean of zero). No information was available for the source position or source instant, therefore, their weights in the regularization matrix were held to be zero. Minimizing Φ in (10) with respect to $\boldsymbol{m}$ results in the regularized solution [12]

$$\boldsymbol{m} = [\mathbf{J}^T\boldsymbol{G}^T\boldsymbol{G}\mathbf{J} + \mu\boldsymbol{H}^T\boldsymbol{H}]^{-1}[\mathbf{J}^T\boldsymbol{G}^T\boldsymbol{G}(\boldsymbol{\delta T} + \mathbf{J}\boldsymbol{m'}) + \mu\boldsymbol{H}^T\boldsymbol{H}\boldsymbol{m'}] \tag{11}$$

The parameter μ is generally chosen so that the χ^2 misfit achieves the expected value of $\chi^2 = M$, for M data. Although it is possible to compute an optimum μ at each iteration of the solution process, μ was held to be constant for this work. Prior to experimentation, simulations revealed that a constant value of $\mu = 100$ would produce sufficient χ^2 misfit and that the final model $\boldsymbol{m}$ would resemble the starting model $\boldsymbol{m'}$ that included the *a priori* estimates of node locations.

The χ^2 misfit of the linear inverse problem (9) was used to derive (11), however, convergence of the iterative process must be determined by the misfit to the nonlinear problem

$$\chi^2 = |\boldsymbol{G(F(m) - T)}|^2 \quad (12)$$

For this work, an iterative solution process was used to find source locations from the RI technique presented above. Convergence of the iteration process was determined by one of two criteria: 1) the change in the *nonlinear* χ^2 statistic between iterations was less than 0.5 percent, or 2) the change in source location between iterations was less than 0.001 m.

Experimental Platform

Hardware

One objective of this work is to produce a system that consists of readily available, commercial off-the-shelf (COTS) products. Accordingly, each of the system's nodes consists of a PC/104 module, a battery pack, and a microphone circuit. The PC/104 module used was a Diamond System's Prometheus that includes data acquisition circuitry, a 100 MHz CPU, 100 Mbps 10/100BaseT Fast Ethernet port, 32 MB of RAM and 128 MB flash disk storage.

Each node's microphone (Panasonic-ECG WM-34BY omni-directional) signal is amplified and fed through a 10 kHz low-pass filter. The resulting signal is then compared to a reference voltage by an electronic comparator whose output is sent to an external interrupt pin on the Prometheus. The external interrupt triggers a subroutine which records the TOA of the acoustic source. After recording the arrival time of the source, group leaders collect arrival time information from all group members and use this data to estimate the position of the source. Both LS and RI position estimates are found within seconds of the source arrival.

A starter pistol is used to generate impulsive acoustic sources that can be heard by all nodes in the network. Measuring the TOA of the source as the point at which the microphone output exceeded the voltage threshold proves to be a robust means of detection for this work because of the impulsive nature of the acoustic wave front. Additionally, the time delay introduced by signal conditioning circuitry is equal for all nodes (within a few microseconds), and therefore is subtracted out in the TDOA calculation. Miniature low-powered ICs have recently been proposed for bearing estimation [14, 15]. With these dedicated hardware components it is possible to measure the TOA to within a few microseconds.

It should be noted that determining the TOA of the source by threshold detection is only effective for environments in which a single acoustic source louder than the ambient noise is present. For environments in which multiple sources are present, a sliding correlator [11], matched filter, or frequency domain method may be used to identify sources of interest, and determine their TOA.

Network nodes communicate via Ethernet ports and a 3Com router. A more practical sensor network would communicate wirelessly, but as it is not necessary for this work, wired communication is used.

Software

Nodes in the distributed system are directed by identical application programs written in C++ that run atop an embedded Linux OS. Network communication is facilitated with use of the Adaptive Communication Environment (ACE), and The ACE ORB (TAO) [16]-[18]. ACE is an open source framework developed for high-performance communication services, and TAO is an open source extension to ACE that arranges the client/server communication in an object-oriented fashion. Using ACE and TAO introduces slight performance reduction to the system, but the benefits of establishing a pattern-oriented structure of program design, while ensuring scalability, robustness and portability outweighs this loss.

Time Synchronization

A distributed acoustic localization system will be accurate only when fine-grained time synchronization between node clocks is available. This is because TOA measurements are made relative to each node's local clock, not a global network reference. Reference Broadcast Synchronization (RBS) [19] provides time synchronization in the sensor network developed for this work. In RBS, an arbitrary node in the sensor network sends a general broadcast packet which is received by other nodes in the network. Receiving nodes mark the arrival time of the broadcast packet on their local clock. The general broadcast arrives at each of the receivers at approximately the same time and can be used as a reference for nodes to compare their clocks. Each node uses this information to setup a table of values representing the difference between their clock and the clocks of other nodes in the network.

Immediately after the reference broadcast is sent, and time synchronization is established, node clocks are synchronized to within a few microseconds. In time however, values of the time difference between node clocks will become inaccurate as the clocks tick at slightly different rates. The average clock drift between any two PC/104 modules was found to be 12 microseconds per second. RBS is performed between five to eight seconds before the source instant during experimentation, resulting in 60-100 microseconds of error in time synchronization (average). A more practical clock skew correction is described in [19] which estimates both the phase offset and the clock skew between nodes.

Grouping

In keeping with the scalable decentralized design objective, the sensor network is divided into subarrays, or sub-groups, of predetermined sizes. The TDOA method of source localization to be solved by LS requires a minimum of four nodes to hear the acoustic source. Thus, a group size of four is the smallest considered for this work. Also considered are group sizes of five, six, seven, eight and nine nodes.

Grouping in this system is performed dynamically and based upon the spacing between nodes. Each node in the system is designated a group leader upon initialization. Leaders then select the n - 1 nodes furthest from themselves to be group members, where n is the group size. Grouping nodes in this manner ensures that groups maintain maximum inter-node spacing, effectively reducing each group's sensitivity to TOA measurement error. This grouping technique is sufficient for the small-scale sensor network developed. For a sensor network containing many nodes monitoring a large area, tradeoffs must be made between inter-node spacing and practical aspects of distributed systems, such as power requirements of wireless communication.

For the scenario in which numerous nodes are spread over a large area, grouping could be accomplished by having nodes choose the *n* - 1 nodes furthest from themselves within some limit. The limit could be in terms of distance (e.g. 20m) or in terms of the number of hops required for two nodes to communicate (e.g. 3 hops). If the network is sparsely distributed, nodes could choose the *n* - 1 nodes nearest to themselves without concern for inter-node spacing.

Experimental Procedure

Three randomly generated arrays consisting of nine nodes and covering a 2.5 by 2.5 meter indoor area were considered in this work. Two of the arrays were used to localize ten near-field sources positioned at random locations within the array area, and five far-field sources positioned outside of the array area. The third array was used to localize five near-field sources. Group sizes of four, five, six, seven, eight, and nine nodes were considered for each of the three random arrays. Group leaders estimated all source positions using both the LS and RI solution techniques previously discussed. The result was 225 near-field and 90 far-field source position estimates for each group size/solution technique pair. One of the three random arrays is plotted in Figure 1 with a set of near-field and far-field source locations. Nodes and sources were placed to within 1 and 4cm of their expected position, respectively. The additional uncertainty in source position was due to the difficulty in holding the starter pistol steady while firing.

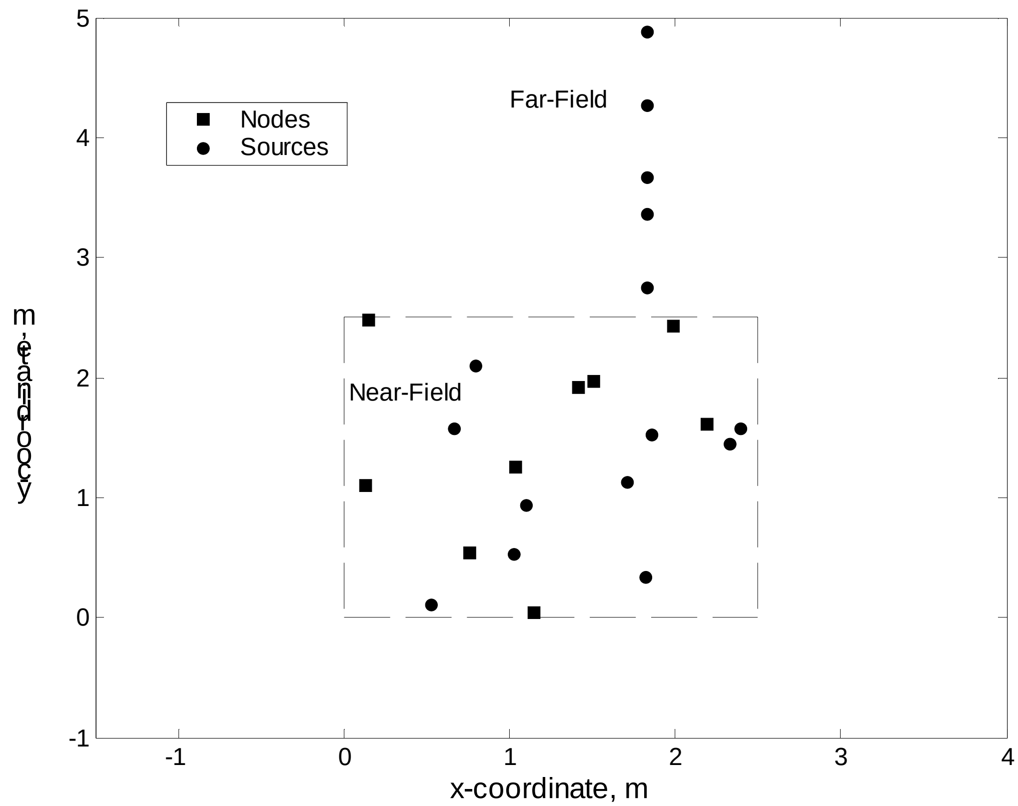

Figure 1. Node and source locations for one random array.

Measurement Error

The ability of the experimental platform to localize sources is limited by TOA measurement error. This error is introduced by clock skew as well as delay in the microphone circuit, threshold detection, and software interrupt. Time synchronization error between nodes at the time of firing was 60-100 microseconds, representing 2-3cm in TOA measurement error. Time delay in the microphone circuit, threshold detection, and software interrupt are approximately the same for each node. Therefore, this delay is subtracted out in the TDOA calculation so that the measurement error introduced by time delay is on the order of a few microseconds.

Experimental results are also limited by node and source placement error which was 1 and 4cm, respectively. Considering both TOA measurement error, and the error in experimental procedure, localization accuracy of 6-7cm is as good as can be expected.

Results

The distributed platform developed for this work was successful in locating sources placed in the near-field using both LS and RI decentralized algorithms. Seen in Figure 2 is the average root-mean-square (RMS) error in source position estimate as a function of group size. This plot shows that the accuracy in source position estimate improves dramatically as the group size is increased, especially for the case of LS. When the system in (3) is exactly determined (4 nodes in each group), the average error in LS position estimate is 32cm. By adding two nodes to each group, the average error in source position estimate is reduced by 20cm.

The RI method of localization found accurate source positions for both small and large group sizes, and generally outperformed LS. The average RMS error in RI source position estimate for a group size of four is 12cm, a significant improvement over the LS method. Increasing the group size improved source position estimates found by RI until a minimum average error of 6cm was reached for a group size of nine nodes. These results are similar to those found in the literature. Chen reported an RMS localization error of about 7 cm for sources in the near field [5].

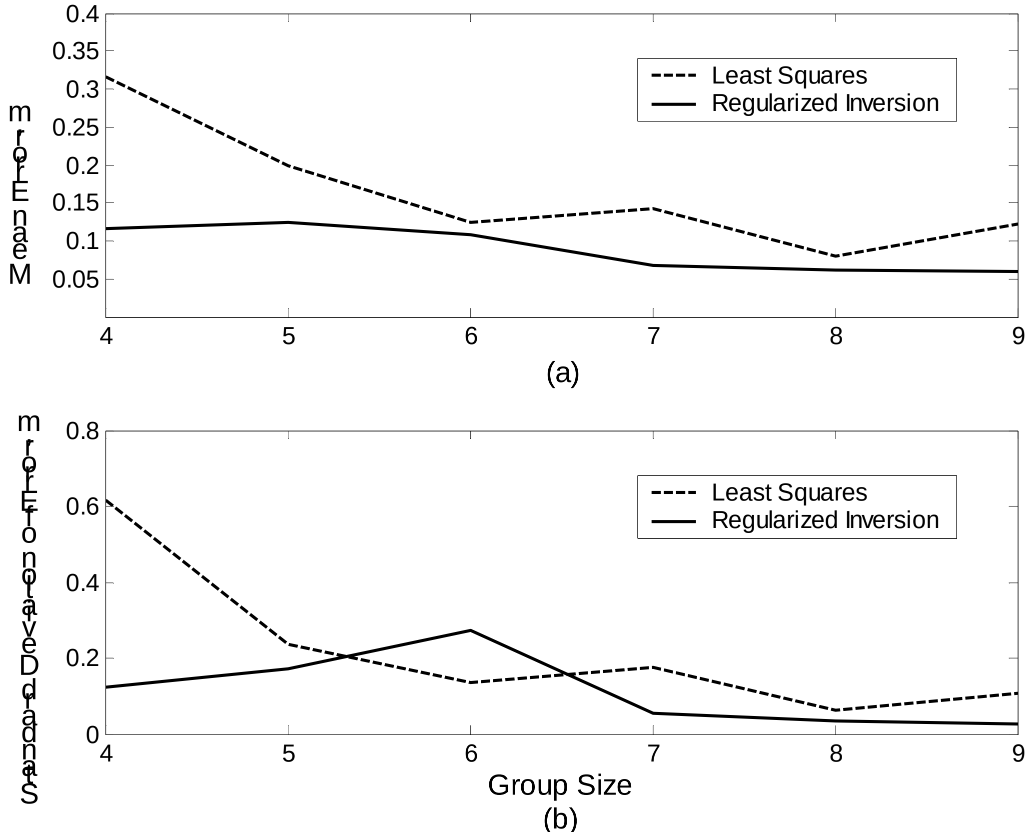

Figure 2. The (a) average and (b) standard deviation of the RMS error in source position estimate as a function of group size. This includes all sources inside of the array area for the three arrays considered.

Error distribution for the LS and RI solution methods can be seen in Figures 3 and 4. It is important to note that the large majority of position estimates found by RI and LS (group size of six or more) are accurate to within 15cm. For example, 96% of position estimates found by RI for a group size of seven are accurate to within 15cm. This is a desirable result when considering the decentralized nature of the localization algorithm, because estimated source positions do not have to be validated by other groups in the network, removing additional communication demands from the system. Outliers seen in Figure 3a (exactly determined LS) occur when the system matrix in (3) is ill-conditioned.

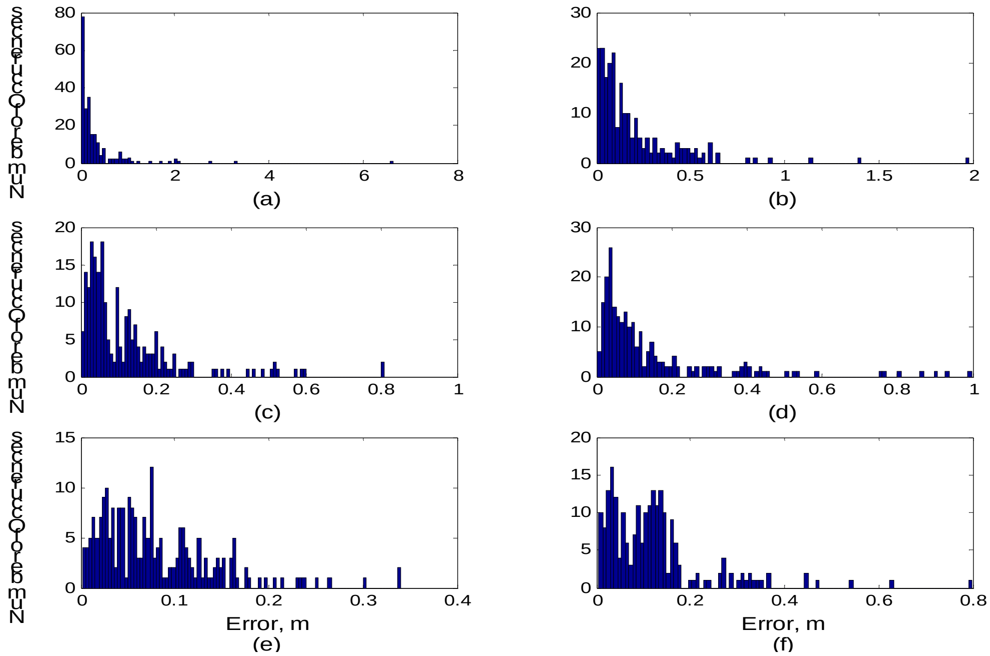

Figure 3. Distribution of the error in source position estimate found by LS for group sizes of (a) four, (b) five, (c) six, (d) seven, (e) eight, and (f) nine [Note: Axes possess different scales].

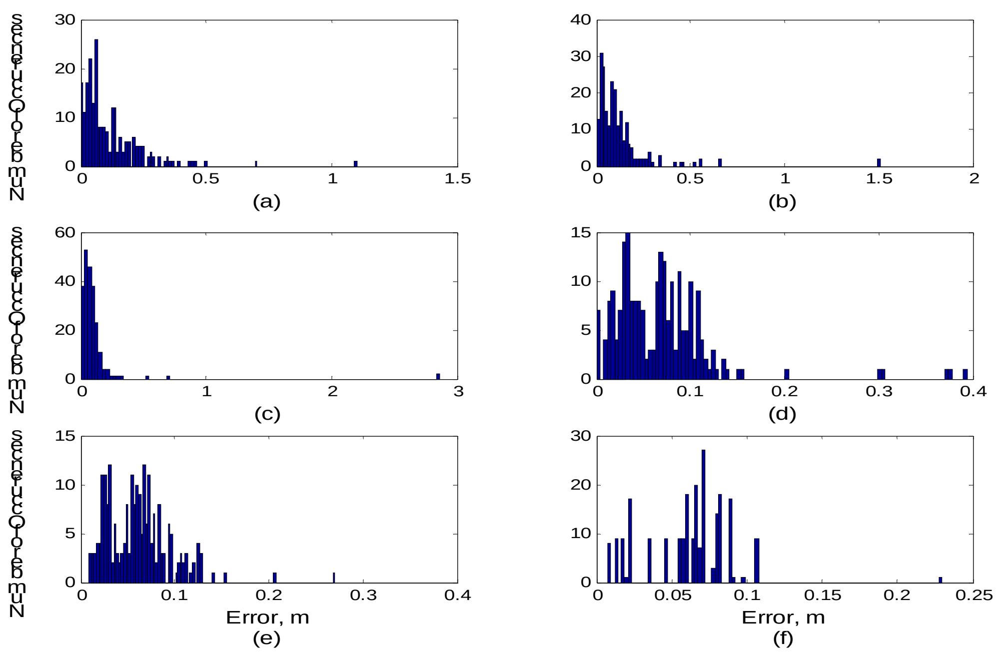

Figure 4. Distribution of the error in source position estimate found by RI for group sizes of (a) four, (b) five, (c) six, (d) seven, (e) eight, and (f) nine [Note: Axes possess different scales].

The distributed platform was not successful in locating sources in the far field. Figure 5 shows the average RMS error in far-field source estimates as a function of group size. The far-field source estimates improve slightly as the group size is increased, but do not reach the accuracy level of the near-field source estimates. Sources in the far-field may be better localized using other methods, e.g. beamforming, and DOA. [5, 6, 8].

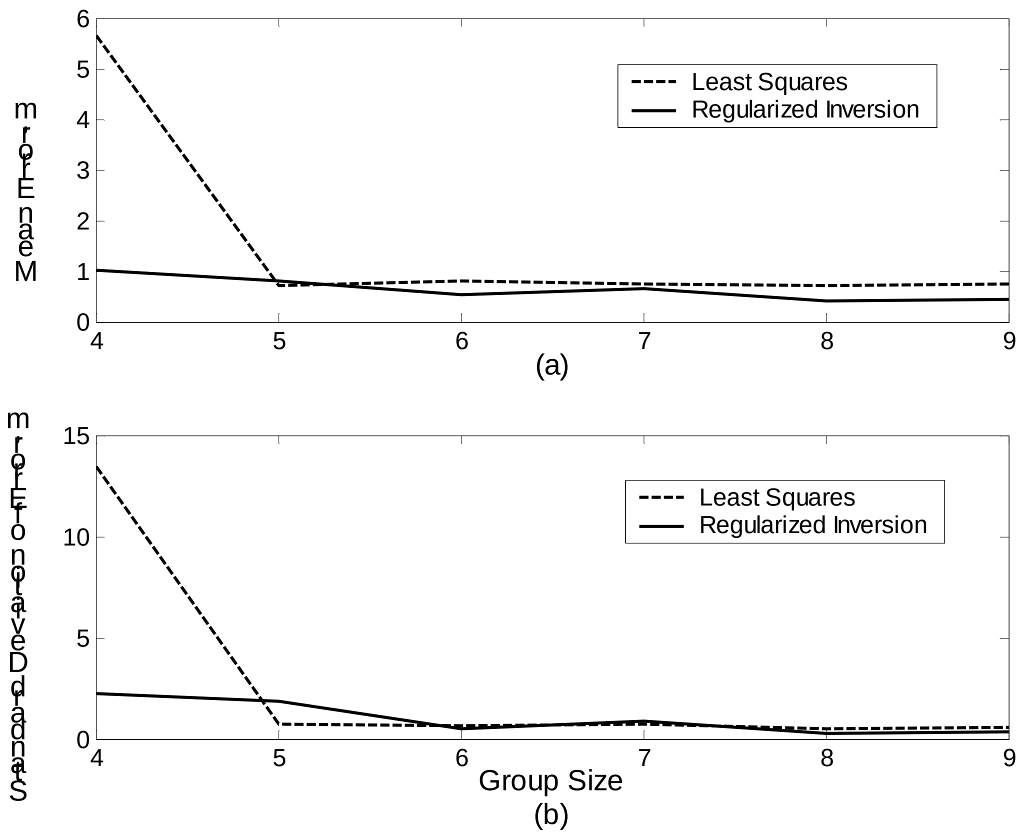

Figure 5. The (a) average and (b) standard deviation of the RMS error in far-field source position estimate as a function of group size. This includes all sources outside of the array area for the three arrays considered.

Conclusions

A distributed sensor network capable of locating near-field sources via decentralized localization algorithms has been presented. Nodes may be added to the network without increasing the complexity of the system, or the computational workload of a central controller. Experimental results confirm that the LS method of localization performs poorly when the system is determined or slightly overdetermined. Increasing the group size of the subarrays improves the performance of the LS method by reducing the likelihood that the system matrix in (3) is ill-conditioned. The RI technique for localization was formulated with use of the regularization method and does not fall victim to the ill-posed nature of the inverse problem as LS does. Accordingly, the RI method provided accurate source position estimates for small group sizes.

Future work includes tracking a moving source in the near-field. The experimental setup is also being used to study the effectiveness of decentralized self-localization.

CHAPTER III

SELF-LOCALIZATION

Many applications of distributed sensor networks require network nodes to have a good understanding of their location and the location of other nodes in the network. With knowledge of node locations the network can conserve power by allowing nodes to communicate only with nodes near themselves. The network may also use node locations to determine the geographical location of events, and to track targets moving through the sensor field. Node locations may be known prior to system use in planned networks, but for networks that are deployed in an ad-hoc fashion, self-localization may be the only method available for determining network topology.

In chapter II a distributed sensor network capable of localizing acoustic sources to within 10cm was presented. Nodes in the network were placed and measured by hand prior to experimentation, and the average error in node placement was about 1cm. The ability of this system to localize sources will deteriorate when node locations are known with less certainty. It is therefore necessary that node positions determined by self-localization be accurate to within a few centimeters (to achieve similar performance).

The following section presents two methods for self-localization. Each is based on TDOA and is implemented in the sensor network in a fully distributed manner. Experimental results are then presented which consider the effectiveness of each method.

Theory

Determining node locations from TOA measurements will again represent an ill-posed, nonunique, nonlinear inverse problem. Both LS and RI solution techniques are considered here. In the previous formulations for source localization, source coordinates (u,v) were determined from knowledge of node locations, the speed of sound, and TOA measurements of the acoustic source at each node. For the process of self-localization, node coordinates (x,y) will be determined from knowledge of the speed of sound, the TOAs of multiple sources at each node, and the locations of the acoustic sources.

The TDOA method presented in [13] is again used to formulate the self-localization problem. If we define the TOA of the j^{th} acoustic source at the i^{th} node as T_{ij}, the time-difference-of-arrival (TDOA) between a reference node and any other node is

$$T_{1ij} = T_{ij} - T_{1j} \tag{13}$$

where T_{1j} is the absolute arrival time of the j^{th} acoustic source at node 1. The locations of reference nodes may be found prior to self-localization by GPS or some other method. The two-dimensional set of nonlinear equations representing the distances between S sources and N nodes can be written as

$$\begin{gathered} d_1^{\ 2} = (x_1 - u_1)^2 + (y_1 - v_1)^2 \\ (d_1 + cT_{121})^2 = (x_2 - u_1)^2 + (y_2 - v_1)^2 \\ \vdots \\ (d_1 + cT_{1N1})^2 = (x_N - u_1)^2 + (y_N - v_1)^2 \\ \vdots \\ d_S^{\ 2} = (x_1 - u_S)^2 + (y_1 - v_S)^2 \\ (d_S + cT_{1NS})^2 = (x_N - u_S)^2 + (y_N - v_S)^2 \end{gathered} \tag{14}$$

where d_j represents the distance from the reference node to the j^{th} source (u_j, v_j), (x_i, y_i) is the position of the i^{th} node, and c is the speed of sound. Expanding these equations and then substituting equations for d_j^2 into the remaining equations results in the following linearized system of $S(N\text{-}1)$ equations in the form

$$\boldsymbol{Ax} = \boldsymbol{b} \tag{15}$$

$$\boldsymbol{A} = \begin{bmatrix} -2u_1 & -2v_1 & 0 & 0 & \cdots & 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & -2u_1 & -2v_1 & \cdots & 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & -2u_1 & -2v_1 & 0 & 0 & \cdots & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -2u_S & -2v_S & 0 & 0 & \cdots & 0 & 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & -2u_S & -2v_S & 0 & 0 & \cdots & 1 \end{bmatrix} \tag{16}$$

$$\boldsymbol{b} = \begin{bmatrix} B_1 + 2cT_{121} + c^2 T_{121}{}^2 \\ B_1 + 2cT_{131} + c^2 T_{131}{}^2 \\ \vdots \\ B_1 + 2cT_{1N1} + c^2 T_{1N1}{}^2 \\ \vdots \\ B_S + 2cT_{12S} + c^2 T_{12S}{}^2 \\ \vdots \\ B_S + 2cT_{1NS} + c^2 T_{1NS}{}^2 \end{bmatrix} \tag{17}$$

$$\boldsymbol{x} = \begin{bmatrix} x_2 & y_2 & x_3 & y_3 & \cdots & x_N & y_N & R_2 & R_3 & \cdots & R_N \end{bmatrix}^T \tag{18}$$

where

$$R_i = x_i{}^2 + y_i{}^2$$

$$B_j = R_1 - 2u_j x_1 - 2v_j y_1 \tag{19}$$

This system of equations contains $3(N\text{-}1)$ unknowns and may be solved by LS when three or more sources are present.

In chapter II, an iterative linearized inversion technique was presented for source localization. This technique was adapted from Dosso's formulation for self-localizing horizontal line array elements [12]. Equations for the TOAs of a single source at N nodes were presented in (4) and may be written for the TOAs of S sources at N nodes. This set of SN nonlinear equations can be linearized and solved using the same iterative solution process presented in chapter II. The model vector $\boldsymbol{m}$ from (5) will be $\boldsymbol{m} = [x_1, y_1, \ldots, x_N, y_N, u_1, v_1, \ldots, u_S, v_S, \gamma_1, \ldots, \gamma_S]$ for this formulation. Source locations and the location of the reference node will be included as *a priori* information here. The regularization matrix in (10) is $\boldsymbol{H} = \mathrm{diag}[1/\sigma_{x1}, 1/\sigma_{y1}, 0, 0, \ldots, 1/\sigma_{u1}, 1/\sigma_{v1}, \ldots, 1/\sigma_{uS}, 1/\sigma_{vS}, 0, 0, \ldots, 0, 0]$. Where σ_{x1}, and σ_{y1} are the standard deviations of the x and y-coordinates of the reference node, and σ_{uj}, σ_{vj} are the standard deviations of the x and y-coordinates of the sources (again assuming that variables are Gaussian-distributed with mean of zero). $\boldsymbol{H}$-matrix weighting of unknown node locations and source instants are held to be zero because no *a priori* information is available for them.

Node locations are found from the iterative solution process presented in Chapter II with the adjustments presented above. Convergence of the iterative process is decided by one of two criteria: 1) the change in the *nonlinear* χ^2 statistic between iterations is less than 0.5 percent, or 2) the changes in node locations between iterations are less than 0.001 m.

The iterative solution process will at times converge to a solution which does not accurately reproduce the network topology. When this occurs, the nonlinear χ^2 statistic does not achieve the expected value of $\chi^2 = M$ for M data. Inaccurate solutions to the iterative process can be thrown out when the nonlinear χ^2 statistic does not achieve its expected value.

Experimental Procedure

The experimental platform described in Chapter 2 is again used to explore the effectiveness of LS and RI self-localization techniques. LS and RI algorithms are implemented using the decentralized grouping approach described previously. Trials were conducted for three randomly distributed arrays of nine nodes placed over a 2.5 by 2.5 meter indoor area. Group sizes of four, five, six, seven, eight, and nine nodes were considered for each of the three random arrays. Self-localization was performed for each group size after groups collected data from five, six, seven, eight, nine, and ten acoustic sources. Approximately 10 % of RI solutions were thrown out based on the χ^2 statistic.

Results

The average RMS error in node location as a function of group size can be seen in Figure 6 for both LS and RI self-localization. The RI algorithm performed slightly better than the LS, but both found node locations to within 6cm (average). Increasing the group size from four to nine nodes reduced the average error in node position by about one centimeter for the RI case.

The average RMS error in node location as a function of number of sources is seen in Figure 7. Node location estimates improve slightly as sources were added for both LS and RI methods. It should be noted that self-localization can be performed when only three or four sources are present. For these cases, the system of equations in (3) is determined, or slightly overdetermined, and the likelihood of the system being ill-conditioned is greater than when five or more sources are present. Although it is not experimentally validated, LS performance will likely suffer when fewer than five sources are present.

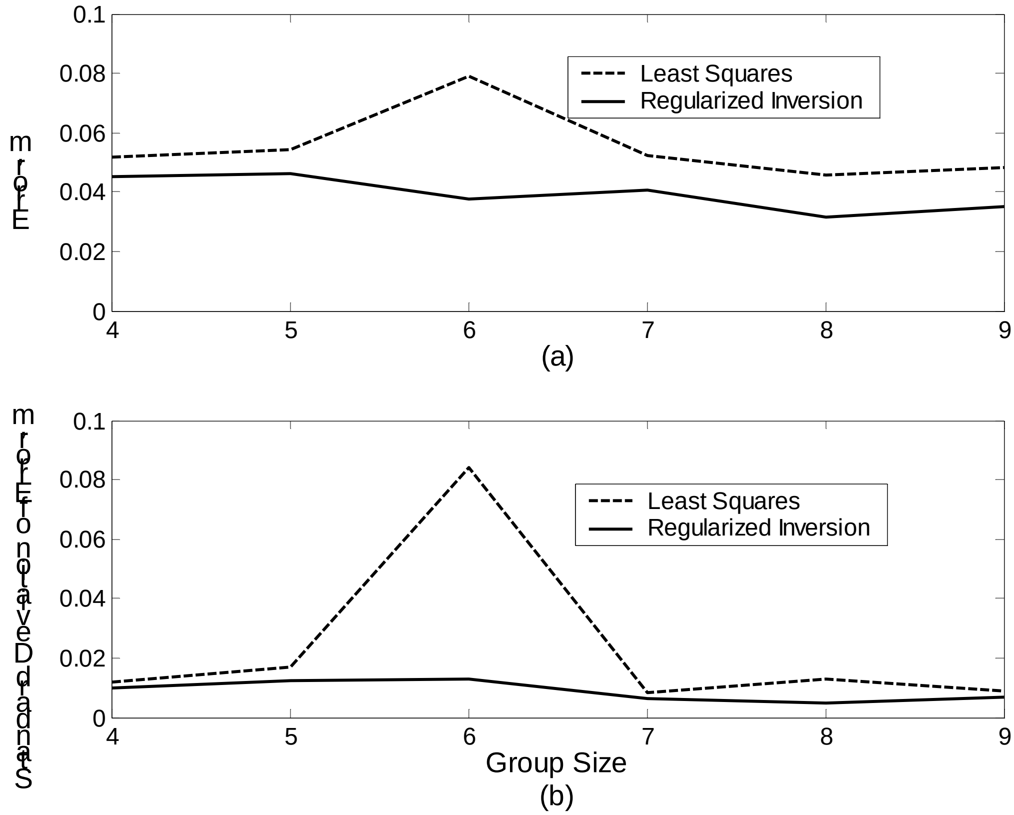

Figure 6. The (a) average and (b) standard deviation of the RMS error in node position estimate as a function of group size.

The self-localization approach taken by Estrin and Girod, in which each node emits an acoustic signal to determine ranging, resulted in error of about 11cm [11]. This is comparable to the 6cm error found in this work. It is therefore feasible to use TOA information of acoustic signals emitted by sources unattached to the network when performing self-localization.

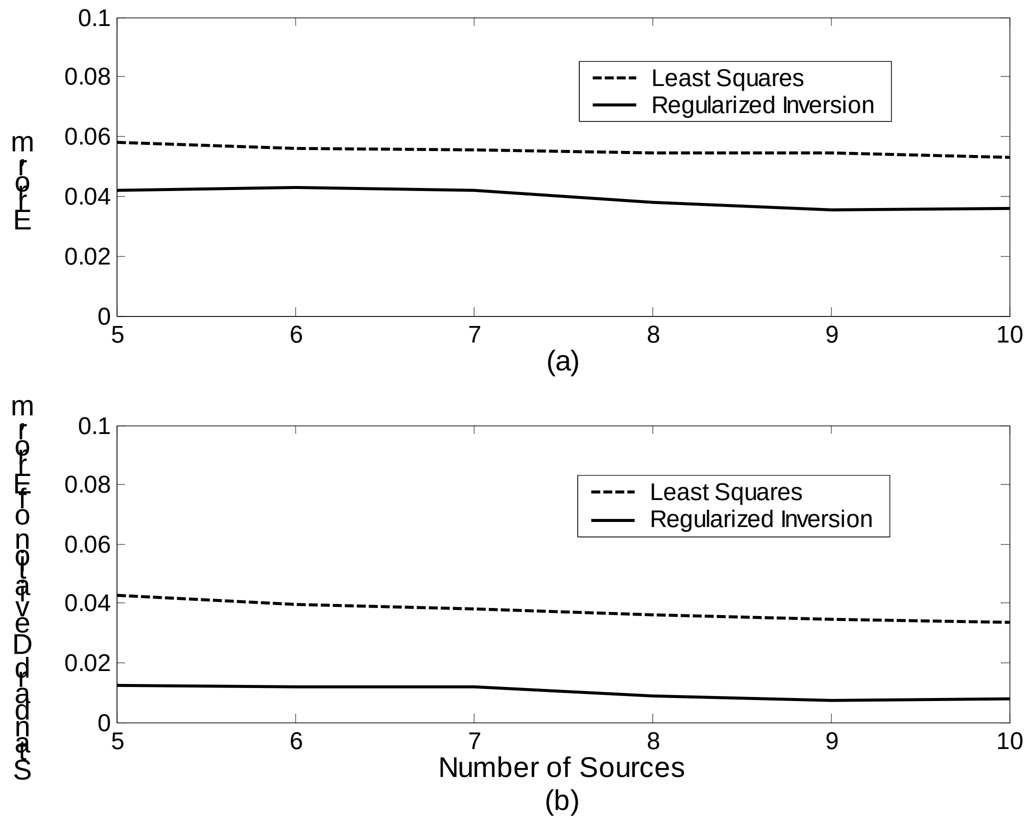

Figure 7. The (a) average and (b) standard deviation of the RMS error in node position estimate as a function of the number of sources

Conclusions

A distributed sensor network consisting of COTS products has been presented in which nodes may be added to the network without increasing the complexity of the system, or the computational workload of a central controller. Experimental results support the use of decentralized algorithms for source and self-localization.

When implementing a decentralized algorithm for source and self-localization engineers must consider tradeoffs between system performance and practical aspects of distributed networks. Localization results may be improved by solving TDOA equations by RI at the cost of higher computational demands. Engineers should therefore select a solution method which provides adequate accuracy without placing significant computational burden on the network. The size of a system's subarrays, or groups, is also an important parameter to consider when performing source localization. An optimum group size should be chosen to minimize communication without sacrificing significant performance.

REFERENCES

[1] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. "Next Century Challenges: Scalable Coordination in Sensor Networks," In Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking, Seattle, WA, August, 1999.

[2] C. Chong and S.P. Kumar, "Sensor networks: evolution, opportunities, and challenges," IEEE Proceedings, Vol. 98, No. 8, pp. 1247-1256, 2003.

[3] N. A. Lynch, Distributed Algorithms, Morgan Kaufmann, 1996.

[4] H. L. Van Trees, Optimum Array Processing, Wiley, New York, 2002.

[5] J.C. Chen et al., "Coherent Acoustic Array Processing and Localization on Wireless Sensor Networks," Proceedings of IEEE, Vol. 91, No. 8, pp. 1154-1162, 2003.

[6] K. Yao et al., "Blind Beamforming on a Randomly Distributed Sensor Array System," IEEE Journal on Selected Areas in Communications, Vol. 16, No. 8, 1998.

[7] R. Brooks, P. Ramanathan, A. Sayeed, "Distributed Target Classification and Tracking in Sensor Networks," Proceedings of the IEEE, Vol. 91, No. 8, 2003.

[8] A. Lendeczi, et al., "Multiple Simultaneous Acoustic Source Location in Urban Terrain," IPSN 05, CD-ROM, Los Angeles, CA, April, 2005.

[9] J. C. Chen, K. Yao, R. E. Hudson, "Source Localization and Beamforming," IEEE Signal Processing Magazine, March, 2002.

[10] K. D. Frampton, "Acoustic Self-Localization in a Distributed Sensor Network," accepted by the IEEE Sensors Journal, March, 2005.

[11] L. Girod, V. Bychkovskiy, J. Elson, D. Estrin, "Locating tiny sensors in time and space: A case study," In Proceedings of the International Conference on Computer Design (ICCD 2002), Freiburg, Germany. September 2002.

[12] S. E. Dosso, M. R. Fallat, B. J. Sotirin, and J. L. Newton, "Array element localization for horizontal arrays via Occam's inversion," Journal of the Acoustical Society of America, Vol. 104, No. 2, pp. 846-859, 1998.

[13] A. Mahajan and M. Walworth, "3-D Position Sensing Using the Differences in the Time-of-Flights from a Wave Source to Various Receivers," IEEE Transactions on Robotics and Automation, Vol. 17, No. 1, 2001.

[14] M. Stanacevic and G. Cauwenberghs, "Micropower mixed-signal acoustic localizer," Proceedings of the 29th European Solid-State Circuits Conference, ESSCIRC '03, pp. 69-72, Sept. 2003.

[15] P. Julian, A. Andreou, P. Mandolesi and D. Goldberg, "A low-power CMOS integrated circuit for bearing estimation," ISCAS '03. Proceedings of the 2003 International Symposium on Circuits and Systems, Vol. 5, pp. V-305 - V-308, May, 2003.

[16] D. C. Schmidt, and S. D. Huston, C++ Network Programming, Volume 1: Mastering Complexity with ACE and Patterns. Addison-Wesley, Boston MA, 2002.

[17] D. C. Schmidt, and S. D. Huston, C++ Network Programming, Volume 1: Systematic Reuse with ACE and Frameworks. Addison-Wesley, Boston MA, 2002.

[18] D. C. Schmidt, D. L. Levine, S. Mungee, "The Design and Performance of Real-Time Object Request Brokers," Computer Communications, Vol. 21, pp. 294-324, 1998.

[19] J. Elson, L. Girod, and D. Estrin, "Fine Grained Network Time Synchronization using Reference Broadcasts," Proceedings of the Fifth Symposium on Operating Systems Design and Implementation, Boston, MA, December 2003.