An Investigation of Localizing Mica2 Motes using the Acoustic ENSBox Platform to Enable a Heterogeneous Sensing Network

Tingyu Thomas Lin
Last Revision: June 8th, 2007

Abstract

This paper presents the findings of an investigation of a heterogeneous acoustic sensing network comprised of the Acoustic Embedded Networked Sensing Box (ENSBox) and the Berkeley Mica2 platforms. The ENSBox platform currently has implemented the key facilities needed for acoustic source localization: a sensor sub-array, wireless network communication, time synchronization, and self-calibration of ENSBox position and orientation. The system expands on the ENSBox platform by including Mica2 motes in performing source localization. The motes, in conjunction with the ENSBox, also can self-calibrate node positions and also has the other necessary facilities needed for acoustic source localization. The ENSBox platform can self-calibrate with about a 2 cm 2D positional error and orientation error of 1˚ in an 80x50 m field (1). From this, the system with at least six ENSBoxes can self-calibrate mote positions with about a 40±23 cm error in ideal simulated environments in a 60x60m field.

1. Introduction

The ENSBox platform is a self-calibrating distributed acoustic sensing platform that performs acoustic source localization. Acoustic source localization is the process of determining the position of acoustic sources. This has many practical applications, ranging from scientific purposes like tracking animal calls (2) to military applications like pinpointing the location of gunshots (3). The ENSBox platform has several facilities that enable acoustic source localization: a sensor sub-array, wireless network communication, and time synchronization. Also, the ENSBox platform can self-localize its own sensor array with a high degree of accuracy. In distributed acoustic sensing networks, it is highly beneficial for sensor arrays to be able to self-calibrate accurately. It is impractical in a large network to manually determine the location of each node, and high accuracy is desirable as error in the position of a sub-array will likely translate to errors in source localization. Each individual ENSBox has a decent amount of resources for computation and a Linux based development environment, making the ENSBox platform ideal in many ways for acoustic sensing applications. In fact, the ENSBox has been used successfully to localize marmots based using marmots alarm calls (2).

However, while trying to localize acoustic sources, the ENSBox platform frequently encounters ambiguity in the recorded data. This ambiguity in collected information may very well translate into errors in acoustic source localization. One possible solution to this problem is to place more ENSBoxes in the field. These added nodes provide additional information to the network, reducing the ambiguity in the aggregate data. However, deploying many ENSBox nodes can be prohibitively costly, both in the sense that each ENSBox is monetarily costly and that to reduce ambiguity the complete functionality of an ENSBox is not needed.

For this purpose, the Mica2 mote platform can be used instead. Mica2 motes as an acoustic sensing platform compares poorly to ENSBoxes: the motes are heavily resource constrained and have relatively low resolution in virtually all measurements they make. Nonetheless, Mica2 motes have the potential to augment the acoustic sensing capabilities of the ENSBox platform. While not being able to provide acoustic localization to the degree of accuracy that the ENSBox can, Mica2 motes can sufficiently act as event detectors for the ENSBox network, and the data collected at the motes' end can be used to improve the accuracy of source localization. The Mica2 motes are also more cost effective to deploy in large numbers, providing a much denser sensing network. And when used in conjunction of with ENSBoxes, the resource constrained motes can offload computations to the ENSBoxes, enabling more complex computations to be performed that a pure Mica2 mote network otherwise couldn't do.

But if the Mica2 motes and ENSBoxes are to act as a single system, accurate mote localization is needed. Accurate localization of the motes is needed to sufficiently relate event detections at motes to the locations the events were detected at. To this end, this paper investigates using the ENSBox platform's acoustic localization facilities, using a direction-of-arrival based localization scheme, to localize speaker-equipped Mica2 motes.

The following sections describe the system and provide an analysis on the system's mote localization accuracy based on simulations that model the system. The simulations are restricted to a 2D plane, though in theory the system is capable of self-localizing and source-localizing in three dimensions. Section 2 of this paper describes the hardware that comprises this system. Section 3 describes integrating Mica2 motes and ENSBoxes into a single system. Section 4 details how the system was simulated in MATLAB. Section 5 describes the experiments and simulation results of localizing the motes. Finally, Section 5 discusses the simulation results and describes what future work can be done.

2. Hardware

The hardware used in the system is comprised of the Acoustic Embedded Networked Sensing Box (ENSBox) and the Mica2 mote.

2.1. Acoustic Embedded Networked Sensing Box (ENSBox)

The ENSBox is built on the Sensoria Slauson board, which has a 400 MHz Intel PXA255 processor, 64MB of RAM, 32MB of on-board flash memory, and PCMCIA slots for a 802.11 wireless card and a Digigram VXPocket440 four-channel sampling card (1). The ENSBox runs a modified version of Linux to support applications run on the ENSBox platform.

Each ENSBox is equipped with a four-microphone sub-array. The four microphones are arranged in an 8 cm square. Two of the microphones at opposing corners are raised above the square's plane by 8 cm. Since the simulations only consider the two-dimensional case of acoustic source

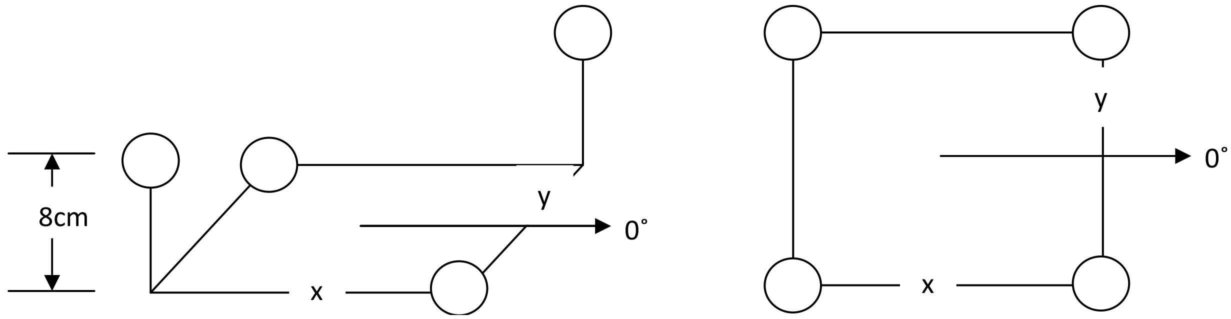

Figure 1. The layout of the microphone array. The diagram on the left is the three dimensional view of the array. However, the simulations only consider two dimensions. The diagram on the right is the two dimensional view that simulations are restricted to.

localization, the fact that two of the microphones are raised is ignored. The hardware is powered by an internal rechargeable battery that allows the system to run for about 24 hours.

The ENSBox already has developed for it the processes needed to self-calibrate positions and orientations. The time synchronization protocol used by the ENSBox requires at least three ENSBoxes to function (see section 3.2. for a description of time synchronization in ENSBoxes), so the smallest deployment of ENSBoxes possible will have three ENSBoxes. A more typical deployment will have 5-10 ENSBoxes to cover a larger field and to tolerate node failures.

2.2. Mica2 Mote

The Mica2 mote consists of a 7.3 MHz 8-bit Atmel ATmega 128L low power microcontroller, a 433 MHz radio, and a port to connect sensor boards to. 4 kB of RAM and 128 kB of flash memory is available for software to run on top of the Mica2 mote platform. To facilitate acoustic self-localization and source localization, a speaker and a microphone needs to be attached as a sensor board. The speaker is used to emit a sound for the ENSBoxes to locate the mote with, and the microphone is used for source localization purposes.

Figure 2. A Mica2 Mote without a sensor board attached.

3. Integrating Mica2 Motes into the ENSBox Platform

Three things are needed to integrate Mica2 Motes into the ENSBox platform:

- Wireless Communication between the two networks

- Time synchronization in the mote network
- Self-calibrating Mica2 motes

Wireless communication is of course essential to a wireless distributed sensing network. Time synchronization is necessary for acoustic localization; it is necessary for all nodes in the network, motes and ENSBoxes, to relate their respective event detection times to each other. Self-calibration of the motes is not strictly necessary, but very useful. It is likely that mote deployment numbers will be large and manually determining the location of motes will be very impractical.

3.1. Wireless Communication

Built into the Mica2 mote is a 433 MHz radio which can be used to communicate with other motes, and the Mica2 motes do provide a MAC layer protocol. The issue of reliable packet transfer across the mote network, which may be critical in implementations of source localization, is not a topic of this paper, but there has been work done in designing reliable data transports like for sensing networks PSFQ (4), RMST (5), and DTNLite (6). For the most part though, wireless sensing networks should be designed to be tolerant of loss packets, as losing a few packets should not detract significantly from the total information gathered by a large sensing network. For the purposes of this paper and simulation, mote communication is assumed to be lossless, which any protocol that provides reliable data transfer across a multi-hop mote network can provide. This is a simplifying assumption, as reliable data transport requires overhead that a practical implementation might seek to avoid, which then the system would have a multi-hop loss-tolerant design in the place of a reliable data transport protocol.

Communication between Mica2 motes and ENSBoxes is greatly simplified by the fact that ENSBoxes can interface with a Mica2 mote via serial connection. In addition, there are libraries for an ENSBox to send and receive packets through the tethered mote. This enables the ENSBoxes and motes to coordinate their behavior and to exchange information across the network.

3.2. Time Synchronization

Time synchronization is critical to acoustic source localization. When a mote detects an event, it must be able to know when it detects the event relative to other motes' event detection times. The internal clocks of the devices are not identical and the clocks will both read different times and increment at different speeds. Given that the difference in times of detection at different motes will be incredibly small due to the speed of sound, the accuracy of time synchronization must be in the order of microseconds. Time synchronization is already a part of the ENSBox; the Mica2 mote network will also need it as well.

The time synchronization mechanism that the ENSBox fundamentally uses is the Reference Broadcast Synchronization (RBS) (7). In a simplified view of RBS, a radio broadcast is made and all nodes in range record the time the packet was received. The difference then between each of the receive nodes' reception time is the clock conversion between them. This requires at least three ENSBoxes: one to broadcast the reference message and two to compare the reception times between. Using this mechanism, the ENSBox platform achieves time synchronization on the order of tens of microseconds across multiple hops (1). However, RBS requires message exchanges between all nodes to communicate time synchronization (8). While this is not an issue in ENSBox networks, where there are only 6 to 10 ENSBoxes, exchanging messages between motes in a large mote network will result in a significant amount of radio traffic overhead for time synchronization.

A better alternative is the Flooding Time Synchronization Protocol (FTSP), proposed by Maroti et al. (8), and is the time synchronization mechanism most ideal for the application. FTSP was designed and implemented on the Mica2 platform, and has been extensively tested in a counter-sniper application (8) (3). In a highly simplified view of FTSP, one mote is elected the root, and its clock is considered to be the reference for global time. The rest of the mote network is organized into a hierarchical tree. The root then sends out a packet with the timestamp of the global time of transmission. The packet is received by the motes that are a layer below the root in the hierarchy. These receiving motes compare their local reception time with the global transmission time, and determine the clock conversion between the two. Then, these motes, like the root, flood time-stamped signals to the motes that are one level down in the mote hierarchy. In this approach, time conversion is arrived at with only pair-wise clock comparisons. Maroti et al. implemented FTSP on a 60-node network and achieved an average per-hop error of about 3 microseconds and a maximum error of 14 microseconds. In three microseconds sound moves .1 cm. This is sufficiently accurate to perform source localization.

Tight time synchronization between the ENSBoxes and Mica2 motes can also be achieved through the tethered mote on each ENSBox. The tethered mote can be modified to partake in FTSP and will have a time conversion table between its internal clock to the global mote time. Similarly, the ENSBox the mote is tethered to will have a conversion table between its own clock and the global ENSBox clock. It is possible to develop a time conversion table between the tethered mote's time and the ENSBox time by sending time-stamped messages through the serial connection. The principles similar in adjusting for the clock differences across a network can be applied to the serial connection to form the time conversion.

3.3. Self-Calibration of Mica2 Motes

The ENSBox has the facilities necessary for source localization; as such, it is natural to attempt localizing motes by equipping the motes with speakers and use the ENSBoxes to localize the tone that the motes can emit. Functionally locating the motes is no different than performing acoustic source localization, with the exception that the motes' emitted tone is controllable. Following this paradigm, localizing a mote consists of the following steps:

1. Schedule the mote to emit a signal
2. ENSBoxes localize the signal
3. Assign back the ENSBoxes' localization results as the mote's location

The motes must be scheduled to avoid having two motes sounding off at the same time, which may potentially confuse results. It may be possible to have multiple motes sound off at the same time and differentiate them during the localization process, but this possibility was not looked at for the purposes of this paper. Therefore, the scheduling mechanism only has to guarantee that no two motes will go off at the same time and enough time elapses between turns and any arbitrary order the motes go in will suffice. With the tight time synchronization, it is possible to directly schedule times for each mote to chirp; alternatively, the ENSBoxes can directly notify a mote when it is that mote's turn. The simplest approach would be for one ENSBox elected as the master to compile the list of motes, tell the mote at the top of the list to emit the calibration signal, and after all necessary localization computations are made and the mote is assigned its estimated location, the master ENSBox marches down the list and repeats this procedure.

There are several possible ways to localize a signal. The one chosen to localize motes is a direction-of-arrival (DOA) based approach, developed for the ENSBox by Ali et al. (2). In (2), the application of source localization was to locate marmot calls; however, their DOA-based source localization is general to any sort of signal and thus adaptable to localizing motes. Their approach was adopted for a few reasons. They have MATLAB scripts already developed to perform acoustic localization, allowing for faster simulating and modeling of the system. More importantly, the simplicity of their approach makes it the most natural method to explore. However, Ali et al. did not quantify the accuracy of their approach, as they lacked sufficient means of determining the true location of marmots to compare results with. The localization takes the following steps:

1. Determining the direction of arrival (DOA) the sound arrives at each ENSBox
2. Combine the DOAs at each ENSBox to determine the origin of the sound

The following two sections explain how these two steps are carried out.

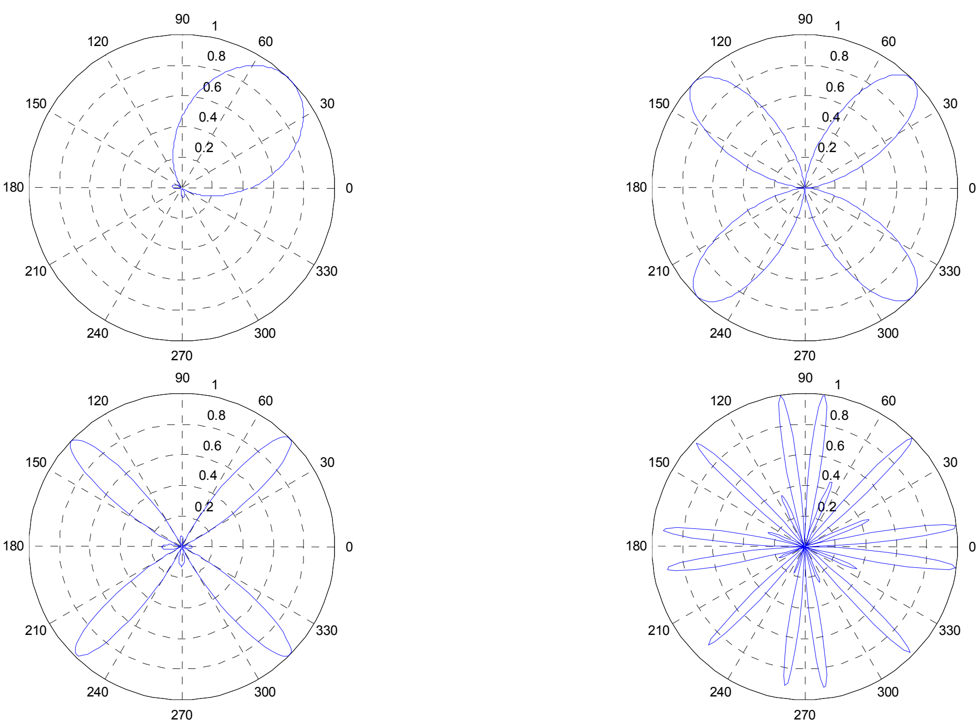

Figure 3. Four DOA likelihood vectors. Each of the four plots is a result of a different frequency signal. Top Left: 1,000 KHz, Top Right: 2,000 KHz, Bottom Left: 4,000 KHz, Bottom Right: 10,000 KHz. The true direction of arrival is 45˚.

3.3.1. DOA-Based Localization

The direction of arrival is determined by the ENSBoxes' four-microphone sub-arrays. A signal propagating through the air will reach each of the four microphones at a different time, resulting in phase differences between the observed signals at each microphone. Since the location of the four microphones is fixed and known, as is the speed of sound in short time intervals, the difference in phases can be used to determine the direction of arrival (DOA) of the sound.

There are inherent ambiguities however in this approach. Determining the DOA estimate relies on the geometry of the sub-array as well as the signal that is received. The result of the DOA estimate is actually done not by determining a single direction value, but a polar plot of likelihood values (see figure 3). There is a limitation on how high a frequency can be. To quote (2):

> "In order to measure the phase of an incoming signal by comparison from two points in space, those two points must lie in the same half-wave. Energy in frequencies with wavelengths shorter than 2x the sensor spacing will be aliased into lower frequencies. This implies that for a sensor spacing of D and signal propagation speed V_s, the maximum frequency detectable without aliasing is $F_c = V_s / (2D)$."

As a consequence, as the frequency of the observed sound increases, the expectation is that the likelihood polar plot will start aliasing, which can be seen in figure 3 as the frequency of call increases. As the signal starts aliasing, grating lobes (aka false lobes, the ones that aren't pointing in the right direction) will start to appear. A solution to this is to move the microphones closer to each other; this will prevent any aliasing. However, this the lobes will widen, introducing more ambiguity into the DOA likelihood plots. Also, the sub-array geometry is fixed in any revision of the ENSBox, making this solution very impractical. As problematic as the grating lobes are, in a multi-node situation where each ENSBox has its own DOA likelihood vectors, the chances of the grating lobes crossing in a way that makes it seem more like the true lobes is highly unlikely. For high frequency calls where the number of grating lobes is large, this rejection will start to fail.

While this concern might be a problem for general source localization, for localizing motes this is very simple to remedy: don't have the motes emit a high frequency signal for localization. Also, given that the highest of call without aliasing is $F_c = V_s / (2D)$, if grating lobes do pose a problem in localization then the motes can call at or below F_c. The microphones are spaced in a 8 cm square, but the largest sensor spacing is in the diagonals of the square, which is 11.3 cm. With the speed of sound at 345 m/s, F_c is about 1500 KHz.

3.3.2. Pseudo-Likelihood Maps

The DOA likelihood vectors that result from a signal mote's call can be fused to produce a most likely position it is at. The DOA vectors are combined in what is called the pseudo-likelihood map (see Figure 4). The search space is divided into 'pixels' and likelihood values are computed for each pixel, based on the value in the DOA likelihood vectors. The pixel with the largest likelihood value at the end is the estimated location of the mote. With this approach, the best-case accuracy is at least the size of the

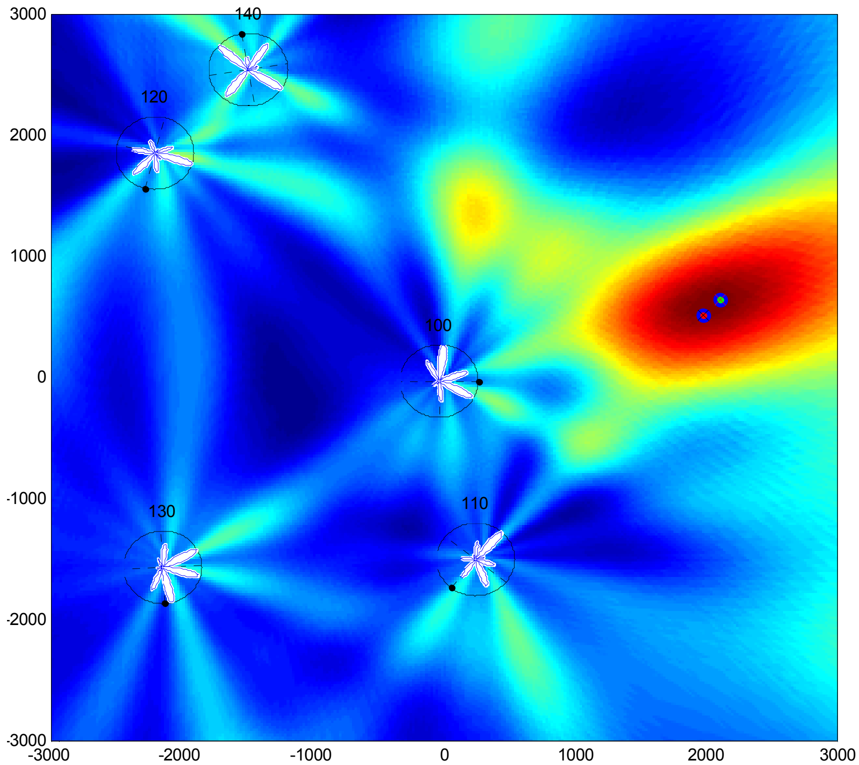

Figure 4. An example pseudo-likelihood map. The five ENSBoxes have produced their own DOA polar plots as a result of the mote's 4 KHz call. The five DOA polar plots are fused together to determine the most likely pixel the mote is in. Each pixel represents a 10x10cm square in the search space. The green dot is the estimated location of the mote while the red dot is the actual location.

pixel. This process is iterated over all the motes in the network to arrive at estimates for mote locations. However, as can be seen in figure 4, the errors in the positions and orientations of ENSBoxes will translate to error in the estimated position of motes, so there will be errors in the self-calibration of the system.

4. Approach to Simulation

The ENSBox and Mica2 acoustic sensing system was modeled and simulated in MATLAB to determine the accuracy of the self-calibration feature. The other two items needed to integrate Mica2 motes into the ENSBox system, wireless communication and time synchronization, are not simulated for this paper even though they are important to acoustic source localization. Time synchronization is not critical for the DOA-based mote localization, and is discussed in this paper for the completeness of an acoustic source localization system. For the purpose of the simulation, wireless communication is assumed to be reliable, mostly as a simplifying assumption but also because there are protocols that provide reliable data transfer. The order of items simulated is:

1. Field
2. ENSBox placement and self-calibration
3. Mote Localization

4.1. Field

The field is modeled as a 60x60 m area with no sound obstructions of any kind. This is a simplifying model, as obstructions such as trees, brush, hilly terrain, etc., pose significant challenges to acoustic localization. The fact that the field is a square is due to the MATLAB function developed by (2) that forms the pseudo-likelihood maps require the search space be a square.

4.2. ENSBox Placement and Self-Calibration

The ENSBox placement is generated randomly in the 60x60 m field, with the constraint that there is a minimum spacing of 10 m between ENSBoxes. This reflects that ENSBoxes are typically spaced out 10 or more meters away from each other in a typical deployment (1), and is reasonable guess as to how a deployment of ENSBoxes will look. To model self-calibration errors, the estimated locations of ENSBoxes are shifted off of the true locations in a random location. The magnitude of the shift is modeled as a Gaussian distribution with a 4 cm standard deviation. This approximates the positional errors reported in empirically observed ENSBox calibration errors (1). The error in orientation error is also modeled after a Gaussian distribution with .96˚, which again is the error based on empirically obtained errors (1).

4.3. Mote Localization

The number of motes and at what time they emit the localization signal does not affect localization results, so localizing motes can be simulated independent of each other. Each mote is randomly generated in the field. The DOA likelihood vectors are simulated by calculating the angle the mote location is to each ENSBox and then generating a set of audio data that the ENSBox would have recorded in such an event. Since the mote can be assumed to be far away from the ENSBox, the sound wave can be modeled as a planar wave, and with a fixed speed of sound the phase offsets in the sound wave recorded at each microphone can be calculated. The sound wave is modeled as a sine wave and does not degrade over distance, though in reality the emitted tone from the mote will not be pure-toned, there will be background noise in the recording, and signal degradation will potentially mean far-away ENSBoxes will not detect the call. The sound is .04 seconds long with a sampling rate of 44 KHz.

The simulated DOA likelihood vectors are used to produce the pseudo-likelihood map. The positions and orientations of the ENSBoxes used to make the map are the estimated positions and orientations that have errors introduced into them. The map is generated with a pixel size of 10x10cm and produces a most likely point the mote is at, and the mote is estimated to be at that location. This estimated location can be compared to the real location to determine the error in the localization process.

5. Experiments and Simulation Results

The accuracy of mote localization was tested along two dimensions: the frequency of the mote's call and the number of ENSBoxes used to localize motes. 20 motes were randomly and uniformly generated in the 60x60 m field. In a real deployment, it may be known that source localization can be improved if more motes are placed in a particular area to gather more information for source

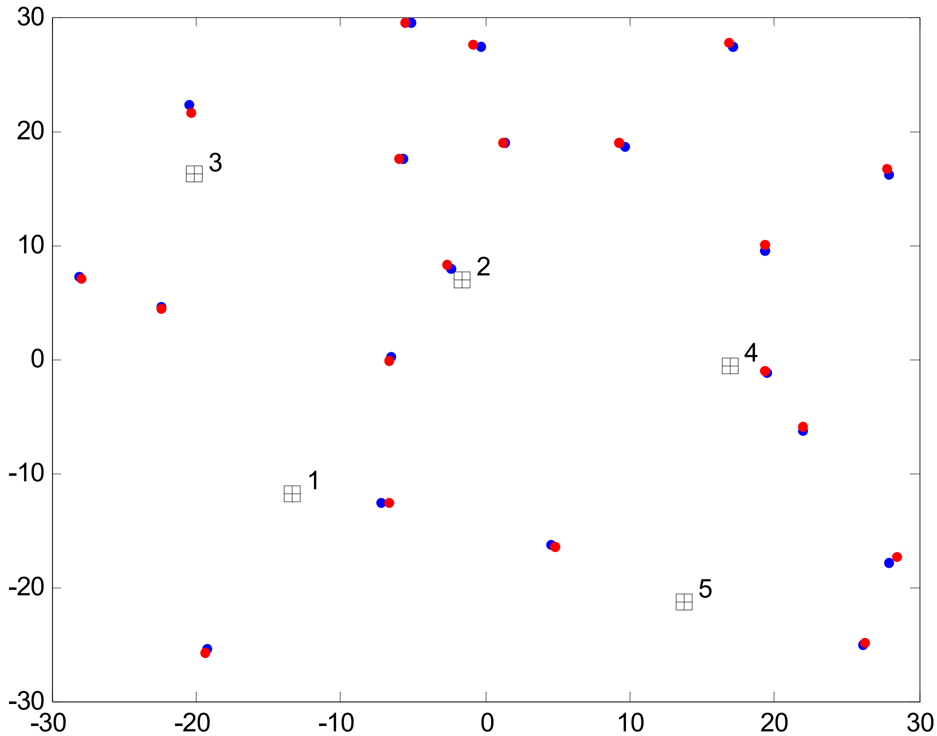

Figure 5. The 20 Mica2 motes and 5 ENSBox positions. The ENSBoxes are the labeled black squares and the Mica2 mote positions are the blue dots. The red dots are the estimated positions of the motes, using 1 KHz calls at the motes.

localization, so the distribution of motes will not be uniform. However, absent of any such prior knowledge, uniformly placing motes in a field a reasonable action. In all experiments, the same 20 mote locations are used.

First experimented with was the frequency of call. Five ENSBoxes were generated in the field with the minimum 10 m spacing constraint, and self-calibration is performed on them once. The errors in positions and orientations are fixed throughout this experiment. The mote localization process was simulated with 1 KHz, 4 KHz, and 10 KHz calls. These three frequencies represent cases with no, some, and a lot of aliasing. This is a test to see if the grating lobes are sufficiently misleading enough to introduce larger localization errors. This was then repeated with different number of ENSBoxes, increasing the ENSBox count up to 8.

The results in table 1 show that in this simulation framework higher frequency calls with more aliasing do not affect accuracy. Simulating the call as a highly idealistic sine wave may lead the pseudo-likelihood maps to have more idealized results. Nonetheless, the principle behind rejecting grating lobes based on the how the lobes intersect is effective in this simulation, even in the extreme cases of 11 grating lobes produced by 10 KHz calls.

Table 1. The mean ± standard deviation errors, in cm. Each cell contains the mean error ± standard deviation in localizing the 20 motes, given the number of ENSBoxes present localizing the motes and the frequency of call. For example, the cell for 5 ENSBoxes and 1 KHz call contains the errors in positions shown in figure 5.

		Frequency of Call		
		1 KHz	4 KHz	10 KHz
Number of ENSBoxes	5	47±23	47±23	47±24
	6	38±23	38±23	39±22
	7	40±35	39±35	36±31
	8	39±25	40±25	38±23

Also, table 1 shows that increasing the number of ENSBoxes from 5 to 8 does not change the errors in localization for this simulation. The average error for the 5-ENSBox configuration seems to be about 7 cm higher, but given the large standard deviations the presence of a 7cm difference in average errors is meaningless. The first experiment was again repeated with 9 and 10 ENSBoxes at only the 1 KHz frequency to see if increasing the number of ENSBoxes even more would change the accuracy. The errors for 9 and 10 ENSBoxes were 44±19 cm and 35±30 cm respectively. This suggests that adding more ENSBoxes beyond 5 does not improve mote localization.

The second experiment is designed to find a lower limit on the number of ENSBoxes before the accuracy starts degrading. The five-ENSBox configuration used in the previous experiment is used again here. ENSBoxes are removed from the five-ENSBox setup and the remaining ENSBoxes are used to localize the motes. All possible combinations of removing ENSBoxes while leaving at least two ENSBoxes were considered. Based on the results in the previous experiment that frequency of call did not change the localization accuracy, this experiment used a 1 KHz call for the motes. Tables 2 and 3 summarize the results in localization.

Table 2: One ENSBox removed, units in cm. The left column is the ID of the ENSBox removed.

1 ENSBox Removed, 4 left	Error in Measurement	
ID	Mean	Std. dev
1	46	18
2	32	15
3	52	24
4	50	24
5	51	32

Table 3: Two and Three ENSBox removed, units in cm.

2 ENSBox Removed, 3 left	Error in Measurement	
IDs	Mean	Std. dev
1,2	52	33
1,3	77	71
1,4	60	22
1,5	86	73
2,3	40	19
2,4	46	40
2,5	33	21
3,4	68	42
3,5	63	32
4,5	62	33

3 ENSBox Removed, 2 left	Error in Measurement	
IDs	Mean	Std. dev
1,2,3	138	114
1,2,4	87	84
1,2,5	73	50
1,3,4	148	157
1,3,5	160	152
1,4,5	133	170
2,3,4	106	257
2,3,5	40	21
2,4,5	88	104
3,4,5	173	242

Leaving only two ENSBoxes for localization produce noticeably worse accuracies. The same is true with three ENSBoxes, though to a somewhat lesser extent. It is worth noting that while the ENSBox platform requires at least three nodes to perform time synchronization, it is possible in a real deployment for ENSBoxes to either fail in sufficient numbers to leave only two ENSBoxes standing. Finally, the case where one ENSBox was removed is somewhat comparable to the original errors from the initial five-ENSBox configuration, though it is still noticeably worse.

6. Conclusion and Future Work

In this simulation framework, the system with at least five ENSBoxes present to localize the motes can self-calibrate the mote positions with an error of 40±23 cm in a 60x60 m field. This result was obtained in an idealized simulated environment with very specific parameters like 20 uniformly placed motes, sine waves as sound wave, etc. This leaves plenty of room for more testing and simulation.

In different deployments, the positions of ENSBoxes will change, as will the positions of the Mica2 motes. However, for the purposed of the experiments the positions of the motes were fixed and the positions of the ENSBoxes are fixed for each number of ENSBoxes. It is currently uncertain how the geometry of the ENSBox placements affects accuracy of mote localization. Different geometries can be simulated and the efficacy of the geometries can be compared to one another.

Noise can also be added to the pure-tone sine wave call the motes emit. In a real deployment, there will be background noise and the signal will degrade as it travels. As such, the ENSBoxes will not simply observe clean sine waves. This will introduce error into determining the DOA likelihood vectors, and the impact of these errors is currently unaccounted for. Signal processing may very well be needed. Several acoustic localization systems (1) (9) (10), the ENSBox included, have signal processing routines to enhance the detection of a signal and to filter out background noise.

Also, a purely DOA approach to localization is inherently sensitive to orientation errors as very small errors in orientation lead to larger errors (11). This limits the effectiveness of the DOA-based localization used to find mote locations. However, the DOA-based localization can be enhanced upon by adding a time difference of arrival (TDoA) component. When a mote emits its localization call, it broadcasts a packet at the same time. The radio packet more or less reaches the ENSBoxes in near instantaneous time, while the sound waves travel at a much slower rate. The time difference of arrival between the packet and the sound at an ENSBox can then be used to determine the distance the mote is away from the ENSBox. This gives additional information to further constrain estimates in mote localization.

Despite the work that remains in improving and analyzing this system, the observed error of 40±23 cm in the simulation results can still be compared to other acoustic localization systems to give context to how well this system performs. To the best of the author's knowledge, the only complete Mica2 mote based acoustic localization platform was developed by Kwon et al. (9). They use a TDoA approach and tested their system with 46 motes in a roughly 60x60 m grassy field. They report a 9.5 m average localization error, but after injecting additional simulated information into the system the average error drops down to 48 cm. No standard deviation is reported in the paper. The simulations achieve accuracy comparable accuracy to Kwon's system. Continued work in investigating improvements like including using TDoA in conjunction with the existing DOA-based localization will potentially lead to better accuracy in mote localization, and continued refinement in simulating and modeling of the system will lead to a better measure of accuracy in localizing the motes.

Acknowledgements

I would like to thank my project advisor Professor Deborah Estrin for giving me the opportunity to work on this research project. I would also like to thank Mike Allen, who has provided the much needed guidance and assistance through the two quarters duration of this project, as well as providing feedback for the drafts of this paper.

References

1. **Girod, Lewis, et al.** The Design and Implementation of a Self-Calibrating Distributed Acoustic Sensing Platform. *SenSys'06.* November 1-3, 2006.

2. **Ali, Andreas M., et al.** An Empirical Study of Collaborative Acoustic Source Localization. *IPSN'07.* April 25-27, 2007.

3. **Simon, Gyula, et al.** Sensor Network-Based Countersniper System. *SenSys'04.* November 3-5, 2004.

4. **Wan, Chieh-Yih, Campbell, Andrew T. and Krishnamurthy, Lakshman.** PSFQ: A Reliable Transport Protocol for Wireless Sensor Networks. *WSNA'02.* September 28, 2002.

5. **Stann, Fred and Jeidemann, John.** RMST: Reliable Data Transport in Sensor Networks. *1st IEEE International Workshop on Sensor Net Protocols and Applications.* May 11, 2003.

6. **Patra, Rabin and Nedevschi, Sergiu.** DTNLite: A Reliable Data Transfer Architecture for Sensor Networks. Fall 2003.

7. **Elson, Jeremy, Girod, Lewis and Estrin, Deborah.** Fine-Grained Network Time Synchronization using Reference Broadcasts. *OSDI 2002.* December 2002.

8. **Maroti, Miklos, et al.** The Flooding Time Synchronization Protocol. *SenSys '04.* November 3-5, 2004.

9. **Kwon, YoungMin, et al.** Resilient Localization for Sensor Networks. *ICDCS.* 2005.

10. **Sallai, Janos, et al.** Acoustic Ranging in Resource-Constrained Sensor Networks. *ICWN '04.* June 2004.

11. **Savvides, Andreas, et al.** An ANalysis of Error Inducing Parameters in Multihop Sensor Node Localization. *IEEE Transactions on Mobile Computing.* November/December 2005.