Thunder: Towards Practical, Zero Cost Acoustic Localization for Outdoor Wireless Sensor Networks

Jingbin Zhang, Ting Yan, John A. Stankovic, Sang H. Son
{jz7q, ty4k, stankovic, son}@cs.virginia.edu
Department of Computer Science, University of Virginia

Abstract--Localization for outdoor wireless sensor networks has been a challenge for real applications. Although many solutions have been proposed, few of them can be used in real applications because of their high cost, low accuracy or infeasibility due to practical issues. In this paper, we propose a practical acoustic localization scheme called Thunder. Thunder employs an asymmetric architecture and shifts most of the complexities and hardware requirements from each node to a single powerful centralized device. The solution is efficient, and requires virtually zero cost in terms of extra per node hardware and in-network communication. This paper also presents an efficient scheduling algorithm called Equilateral Triangle Scheduling to schedule Thunder for very large sensor networks and a resilient algorithm called Adaptive Fuzzy Clustering to provide robust localization without sacrificing efficiency in the presence of a high percentage of large ranging errors. To validate and evaluate Thunder, we built an experimental localization system based on the Mica2 platform, which achieved localization errors of about 1 meter in medium scale localization experiments.

. Introduction

Localization for outdoor wireless sensor networks (WSNs) is a fundamental middleware service for many WSN applications. For example, in military surveillance applications [7] location information of each node is essential to determine a target's position. Although many approaches (e.g., [2][6][9][17][23]) have been proposed to solve the outdoor localization problem, few of them can be used in real applications due to practical issues such as high cost or low accuracy.

Because of the high accuracy of acoustic ranging, several acoustic localization schemes [9][22][23] are proposed for outdoor WSNs. These acoustic localization schemes are mainly based on peer-to-peer acoustic ranging with a certain percentage of anchors and require extra per node devices to perform ranging among neighbors. While these approaches show some promising results, they also have many practical problems. First, for a static WSN, localization only needs to be done once. It is not cost-effective to equip each device for a one-time localization. Second, the effective ranges of these per node devices are constrained by the cost, size or power supply from the nodes. For example, even for the second generation Medusa nodes [23], the peer-to-peer ranging distance is only 10-15 meters. The limited ranging distance places extra requirements on node density in order to get enough range for each node. Third, for a large scale WSN comprised of tens of thousands of nodes, both the cost for a certain percentage of anchors and the cost for in-network communication are large; also, it is very difficult to schedule sound broadcasts at each node to avoid interference from each other to reduce localization time in an efficient and scalable way. Finally, real environments are comprised of various obstacles, such as trees and bushes. The obstacles can cause severe signal attenuation and multi-path signals, resulting in large ranging errors. In the presence of a high percentage of large ranging errors, the distributed localization algorithms used in these localization schemes suffer both from the peer-to-peer ranging errors and error propagation. The end result is that many of these previous solutions don't address large ranging errors or if they do, the performance is poor. All these problems are critical in real applications. In this paper, we present a practical acoustic localization scheme called Thunder for outdoor WSNs. Thunder employs an asymmetric architecture and is able to solve or avoid all these problems effectively with virtually zero cost both in terms of extra per node hardware and in-network communication.

While acoustic based localization is well studied, this paper makes the following five contributions. First, based on a common time difference of arrival (TDOA) technique we present a practical acoustic localization scheme with high accuracy and low cost for outdoor WSNs. Second, we show how to scale the solution to very large networks of 10,000 or more nodes by providing a scheduling algorithm called Equilateral Triangle Scheduling (ETS). Third, we propose an efficient algorithm called Adaptive Fuzzy Clustering (AFC) to provide robust localization for Thunder in the presence of a high percentage of significant ranging errors. AFC can provide accurate localization with average localization errors under 20cm when the percentage of large errors is below 60%. Fourth, we built an experimental Thunder system on the Mica2 platform, which achieved localization errors of about 1 meter, to verify the feasibility of Thunder to support long distance acoustic ranging. Our experimental Thunder system supports effective ranging up to 137 meters. To the best of our knowledge, this is the longest acoustic ranging distance ever achieved in WSNs. Fifth, we identified a hardware saturation problem in the tone detector on the Mica sensor board caused by strong acoustic signals, which makes the tone detector not responsive to acoustic signals. We have developed an efficient solution called Three Phase Adjustment (TPA) to solve this practical implementation issue.

. Thunder System Design

The main idea of Thunder is to use a single centralized device comprised of a speaker, a powerful radio transmitter and a GPS receiver to emulate thunder and lightening in nature that can be heard and seen many kilometers away. With this single centralized device, we only need to move it to 3 different locations not in a line to localize all the nodes in a field in two dimensions.

One practical application scenario we envision is as follows: a helicopter equipped with the powerful centralized device first drops a large number of sensors in a wide area randomly; then the helicopter flies to several locations and at each location, the helicopter simultaneously broadcasts the sound and sends a radio signal containing the location of the helicopter obtained from the GPS receiver; sensors receiving both the radio signal and the acoustic signal can use the standard Time Difference of Arrival (TDOA) approach to compute their distances to the helicopter; after a sensor gets 3 distances to the helicopter at 3 locations not in a line, it can use trilateration to compute its 2D location.

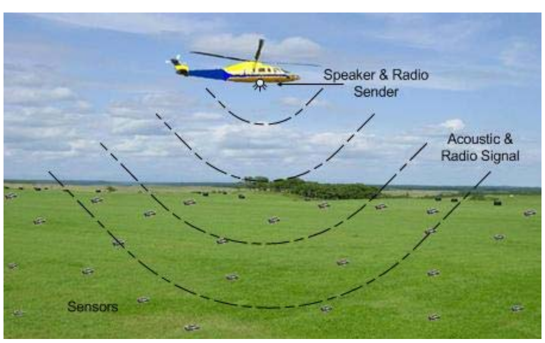

Figure 1: An example of the Thunder localization scheme.

This approach has a number of benefits. First, we move most of the complexities and hardware requirements from each node to a single powerful centralized device. Each node in the field only needs a microphone to capture the acoustic signals and a radio receiver to receive radio messages, which are available to almost all the current popular sensor motes, such as the Mica and XSM series. No extra ranging devices, such as ultrasound transceivers [23] or powerful buzzer units [9], are needed. Further, no anchors are used in the field. Only one GPS receiver is used for the single centralized device. In this way, the cost for extra per node hardware approaches zero. Second, there is no peer-to-peer traffic. This means that there is zero cost in terms of in-network communication. Each node in the field is only responsible for receiving radio and acoustic signals from the centralized device, and then computes its own location independently. Not a single message is generated from the nodes in the field and thus the energy consumption of the nodes for in-network communication is also zero. Third, no assumptions about the topology of WSNs are made. The localization is not affected by the factors such as node density or network connectivity. Fourth, it is efficient and fast to localize a WSN, because we only need to move the centralized device to several locations with a car or a helicopter and at each location broadcast an acoustic signal and a radio signal simultaneously. Fifth, it is scalable as the size of WSNs grows. We design an efficient scheduling algorithm called Equilateral Triangle Scheduling (ETS) described in Section to schedule Thunder for very large WSNs. Finally, it can provide very robust localization without sacrificing efficiency in the presence of a high percentage of large ranging errors by using Adaptive Fuzzy Clustering (AFC), which is described in Section □. Also, it can be easily extended to support 3D localization. For example, a node can use ranging measurements to the centralized device at 4 different locations not in the same plane to compute its 3D location. This single centralized device can be used for many WSNs and thus the cost to build such a centralized device can be amortized to virtually zero.

The efficiency of Thunder mainly depends on the ranging distance a single centralized device can support. In reality, the effective ranging distance of the centralized device is constrained by the power of the speaker and the radio transmitter. Also, it is generally desirable to keep the volume of the sound within human tolerance levels. In order to verify the feasibility of Thunder to support long distance acoustic ranging, we develop an experimental Thunder system on the Mica2 platform, which is shown in Sections □ and □. In our experimental Thunder system, we use a speaker whose maximum sound intensity is 73dB, which is only at the same magnitude level as street noise in a city. In addition, it is important to point out that when we broadcast the sound it is only transmitted for a very short interval of time, i.e., 100 msec. With this sound intensity and time interval we can support effective ranging up to 137 meters. With a more powerful speaker, we expect it to support ranging up to several hundred meters. However when WSNs grow larger, the signals from the centralized device may not be able to cover the whole WSNs. In that case, we still need to move the device to many more than 3 locations to localize the whole area. This scaling issue is discussed in Section .

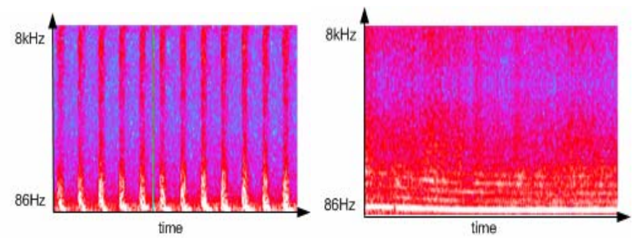

(a) (b)

Figure 2: (a) The spectrum for a typical helicopter. (b) The spectrum for a typical car.

We need to consider whether the helicopter noise or car engine noise interfere with the acoustic signals from the speaker. As shown in Figure 2, most of the sound energy from helicopter and car noises is less than a 1 kHz frequency. By choosing acoustic signals of a higher frequency, the interference from the helicopter and car noises can be avoided.

. Thunder for Very Large Wireless Sensor Networks

When WSNs grow larger, the signals from the centralized device may not be able to cover the whole area. In this section, we describe the scheduling algorithm called Equilateral Triangle Scheduling (ETS) to schedule Thunder for very large WSNs over many square kilometers, and discuss the impact of Doppler effects on acoustic ranging.

.A. Equilateral Triangle Scheduling

When WSNs become large, in order to make sure that every node in the field has enough ranging measurements to the speaker, we need to schedule where to broadcast sound to minimize the number of broadcasts.

We make the following three assumptions in order to simplify the problem. First, we assume that we are using an omni-directional speaker. We define the speaker's effective ranging distance as the distance within which sensors can detect its acoustic signals reliably. Then we assume that the speaker's effective ranging distances to different sensors are the same. So the area in which sensors can detect the sound reliably is a perfect circle. Finally, we do not take the edge effects into account, because we are considering a large area over many square kilometers.

We model the problem in the following way: the speaker's effective ranging distance is z; the sensor field to be localized is a rectangle of size $x \times y$; the question is at which locations should we broadcast the sound to minimize the number of broadcasts, meanwhile satisfying the constraints that every point in the sensor field is covered by at least three different circles of radius z whose centers are the speaker's locations, and the centers of at least three circles are not in a line.

This problem is more complicated than the circle covering problem [28], which is to get the lower bound for a covering using equivalent circles. The circle covering problem only requires each point in the plane to be covered by at least one circle. In [8], Kershner derives the tight lower bound for the number of equivalent circles that covers a geometric area and therefore proves that the total area of the circles is at least $2\sqrt{3}\pi/9$ times the area to be covered. The constant $2\sqrt{3}\pi/9 \approx 1.209$ "may be thought of as measuring the proportion of unavoidable overlapping" [8].

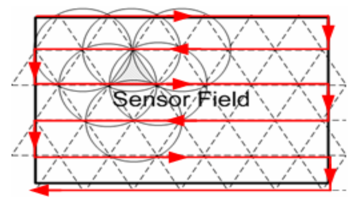

Figure 3: Broadcast scheduling in a very large scale WSN. The arrows show the route to traverse the vertices.

However, our problem requires each point to be covered by at least three different circles and the centers of at least three circles are not in a line, which to the best of our knowledge does not have a tight bound proven in publications. Therefore, we propose an efficient solution called Equilateral Triangle Scheduling (ETS) that provides the same proportion of redundancy $2\sqrt{3}\pi/9$ as in the optimal solution of the circle covering problem. In ETS, we first use equilateral triangles to divide the sensor field as shown in Figure 3. The length of the edge of the equilateral triangles is z. Then we only need to broadcast the sound at the vertices of these equilateral triangles.

As we can see from Figure 3, by using ETS each triangle is fully covered by 3 different circles, so we can make sure that every point in the sensor field is covered by at least three different circles whose centers are not in a line. Also, each triangle is partially covered by another 3 different circles. This is the overhead we incur by using ETS. We get the average number of circles that each point in the sensor field is covered:

$$\left(\pi \times z^2 \times \frac{1}{6} \times 3\right) \Big/ \left(\frac{1}{2} \times z \times \frac{\sqrt{3}}{2} z\right) = 2\sqrt{3}\pi / 3 = 3.627 \quad (1)$$

By dividing Formula (1) by 3, we can show that the proportion of redundancy by using ETS is $2\sqrt{3}\pi/9$, exactly the same as the unavoidable overlapping proportion for the circle covering problem.

Based on ETS, we only need to move the centralized device line by line along the vertices of the equilateral triangles as shown in Figure 3 and broadcast the sound signals and radio signals simultaneously at each vertex to perform localization.

.B. Impact of Doppler Effects

When the centralized device is broadcasting the acoustic and radio signals for ranging, one option is to let the centralized device stop at each scheduled location, and then broadcast. To reduce the time for localization, it is desirable not to stop the centralized device when broadcasting. However, this option is constrained by Doppler effects.

The speed of sound does not change with a moving sound source and a static listener, while the frequency of the sound changes. Because the receivers only listen to a certain frequency range of the sound in order to avoid the interference from the environment, we should control the speed of the moving sound source to make sure that the resulting frequencies are still in the valid listening frequency range of the nodes.

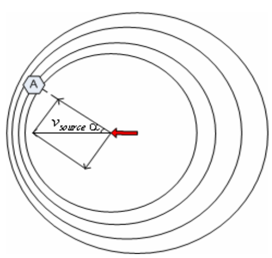

Figure 4: Doppler effects

Figure 4 shows an example of Doppler Effects. The sound source is moving from right to left with the speed of v_{source}. We can compute the frequency of the sound that node A receives as the following:

$$f_{receive} = \frac{v_{sound} \times f_{original}}{v_{sound} - v_{source} \times \cos \alpha}, \quad (2)$$

where $f_{original}$ denotes the original frequency from the sound source, $f_{receive}$ denotes the received frequency, and v_{sound} denotes the speed of sound. Suppose the frequency range that nodes listen to is [f_{min} , f_{max}], in which f_{min} denotes the minimum frequency the nodes can detect, and f_{max} denotes the maximum frequency the nodes can detect. To make sure the resulting frequency is detectable, $f_{receive}$ should be within [f_{min} , f_{max}] as expressed in Formula (3):

$$\frac{v_{sound} \times f_{original}}{v_{sound} + v_{source}} > f_{\min} \ \&\& \ \frac{v_{sound} \times f_{original}}{v_{sound} - v_{source}} < f_{\max} \quad (3)$$

We can get the speed limit of the sound source in Formula (4):

$$v_{source} \le \min\left(\frac{v_{sound} \times (f_{original} - f_{\min})}{f_{\min}}, \frac{v_{sound} \times (f_{\max} - f_{original})}{f_{\max}}\right) \quad (4)$$

Consider our experimental Thunder system described in Section □ as an example. In our experimental Thunder system, $f_{original}$ equals 4.7kHz, and the frequency range is [4.3kHz,5.1kHz]. We use 340m/s as the speed of sound. Based on Formula (4), assuming that the speed of the centralized device does not change, we compute that the maximum speed of the centralized device is 26.7m/s, which is around 96km/hour, a reasonable speed limit for a helicopter or a car.

.C. Localization Time

Based on the route shown in Figure 3, we can estimate the length of the route the centralized device needs to traverse:

$$Length = \min\left(\left(\left\lceil \frac{x}{z \times \sqrt{3}/2} \right\rceil + 1\right) \times y + x, \left(\left\lceil \frac{y}{z \times \sqrt{3}/2} \right\rceil + 1\right) \times x + y\right) \quad (5)$$

From Formula (5), the time for localization can be estimated as follows:

$$t = \frac{\min\left(\left(\left\lceil \frac{x}{z \times \sqrt{3}/2} \right\rceil + 1\right) \times y + x, \left(\left\lceil \frac{y}{z \times \sqrt{3}/2} \right\rceil + 1\right) \times x + y\right)}{\min\left(\frac{v_{sound} \times (f_{original} - f_{\min})}{f_{\min}}, \frac{v_{sound} \times (f_{\max} - f_{original})}{f_{\max}}\right)} \quad (6)$$

Consider the largest WSN assembled to date [29], which covers an area of 1.3km × 300m. If the speaker's effective ranging distance is 200m and the maximum speed is 26.7m/s computed from Section .B, it takes only 157 seconds to localize this large area.

. Improving Robustness of Thunder

In real applications, the environment may be complicated. For example, the area where we deploy the WSN may have many obstacles, such as trees and bushes, causing severe signal attenuation and multi-path signals. The signal attenuation and multi-path signals may result in large ranging errors. However, Thunder provides a feasible way to provide robust localization in the presence of a high percentage of large ranging errors. In Thunder, if it is acceptable to send many sound signals, the centralized device can keep on sending acoustic signals for ranging periodically when it is traversing the route. For example, the moving centralized device can send out signals for ranging every second, if the maximum ranging distance is smaller than the distance an acoustic signal travels in one second to make sure that there are no multiple acoustic signals for ranging in the field at the same time. Thus each node in the field can get many more than 3 ranging measurements. However, the nodes do not know which ranging measurements contain large errors. So the problem becomes how can we use these extensive ranging measurements to correctly compute the location, even though many ranging measurements may contain large errors?

This problem is similar to what is described in [10][11]. The main difference is that [10][11] consider the large ranging errors to be caused by malicious nodes, while we consider them to be caused by the limitations of ranging methods, due to signal attenuation, multi-path signals and other reasons. In other words, malicious nodes may purposely give wrong measurements and may even collude with each other to coordinate wrong measurements. In our case, because there is only a single centralized device and nodes do not exchange data, it is not feasible to maliciously generate colluding measurements. Colluding measurements are defined as compromised measurements that coordinate with each other to push the localization result toward the same wrong location.

In [10], Li et. al. propose using Least Median of Squares (LMS) to eliminate the outliers. LMS works efficiently when the contamination ratio of the measurements is low, but the computation overhead increases very fast as the contamination ratio increases, because it needs to try more subsets to find a good subset without contamination. Also, it is robust for up to 50% of contamination ratio at most, even though these contaminated measurements do not collude to push the localization to the same wrong location. In [11], Liu et al. propose Minimum Mean Square Estimation (MMSE) and Voting-Based Location Estimation to secure the localization. Among these two methods, Voting-Based Location Estimation shows better performance especially when the contamination ratio is high. Voting-Based Location Estimation can be considered as a simplified version of the probabilistic approach proposed in [24].

We propose an efficient algorithm called Adaptive Fuzzy Clustering (AFC) to solve this problem. AFC not only works extremely well when there are no colluding measurements, but also is very robust when there are multiple colluding measurements.

A. Adaptive Fuzzy Clustering

If ranging measurements are not error-free, a location computed from these measurements is inclined to contain errors. We define a possible location of a node as a location computed from a certain number of ranging measurements by a certain method. For example, we can use m ($m \geq 3$) ranging measurements to compute a possible location by multilateration. The computed possible location can be close to the true location of the node if the ranging measurements used to compute that location contain minimal errors. It is also possible that the computed location is a false location that is far away from the true location if the ranging measurements contain large errors. If a node obtains n ranging measurements, the node can compute C_n^m possible locations for itself.

The key idea of AFC contains two phases. The first phase is to obtain all the possible locations for the node with a certain method. The second phase is to use cluster analysis to find an area with the maximum density of possible locations. Then we use the average of all the possible locations in that area as the final location for that node.

Least Square (LS) and Nelder Mead (NM) have been widely used for multilateration. However, it is computationally prohibitive to use these nonlinear regression methods to compute all the possible locations for resource-constrained sensor nodes. For example, if we use 3 ranging measurements to determine a possible location, it needs to use LS or NM C_n^3 times, a nightmare for a sensor mote. Here, we propose an approximate, but efficient way to obtain all the possible locations.

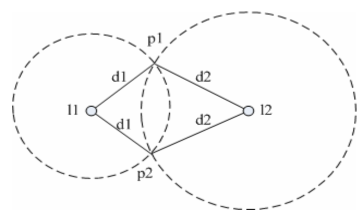

Figure 5: Using two ranging measurements d1 and d2 to get two possible locations p1 and p2.

Phase 1: Our method is based on the following two observations. First, normally two ranging measurements can be used to compute two possible locations for a node. Although which one is correct is unknown, at least one of the two is correct assuming the ranging errors are minimal. As shown in Figure 5, by using two ranging measurements d1 and d2 to locations l1 and l2 correspondingly, the node knows that its location is either p1 or p2. The second observation is that the two possible locations by using two ranging measurements can be computed in linear time and can be very efficiently implemented on motes. In AFC, we compute all possible locations by using any two ranging measurements without caring about which ones are correct, because most wrong locations are eliminated in Phase 2. We initially consider that all the computed locations as possible locations for the node. In this way, we only need to compute the locations with two ranging measurements C_n^2 times, and the computation cost is acceptable.

Phase 2: After we obtain all the possible locations, we use cluster analysis to get the area with the maximum density of locations. The cluster analysis is necessary because locations with large errors might be introduced due to ranging errors and the wrong locations are generated in Phase 1. Because the computation overhead for precise cluster analysis can be prohibitive, we use an approximation to get the area with the maximum density of possible locations. The following briefly describes how the algorithm works.

Step 1: Compute the center of gravity of the set of all the existing possible locations.

Step 2: Compute the average distance l from the set of all the existing possible locations to the center of gravity.

Step 3: Remove all the possible locations whose distances to the center of gravity are larger than $l \times \alpha$ from the set of all the existing possible locations. The value of α is to be determined. From our simulation, we find that a value around 1 is appropriate for α.

Step 4: If the l calculated from Step 2 is below a certain threshold or if the number of rounds exceeds a specified maximum number of rounds, the algorithm terminates and the center of gravity of all the remaining possible locations is used as the location of the node. Else, repeat from Step 1. From our simulations, a value of 10 is appropriate for the maximum number of rounds.

This algorithm works efficiently. The computation complexity of this algorithm is $O(rn^2)$, where r is a constant, whose value is the maximum number of rounds and n is the total number of ranging measurements. However, l converges quickly. Normally it only takes 5 rounds to find the final location.

The main reason that AFC is capable of finding the true location of the node correctly with a minimal error is that if there are enough ranging measurements with small errors, the density of possible locations around the node's true location is higher than other areas. That is because the possible locations with small errors always gather around the true location of the node, while the possible locations with large errors are inclined to spread out if there are no colluding measurements. AFC can also be applied to other localization schemes for which colluding measurements may occur. Using AFC, the possible locations computed from the colluding measurements also concentrate near a false location. However, if the density near the true location is higher than that near the false location, AFC is inclined to converge to the true location as shown in the algorithm description. If the percentage of the colluding measurements is greater than 50%, in which case the density near the false location overweighs that near the true location, no algorithm is able to find the correct location with a high probability unless extra information is provided.

. B. Evaluation of Adaptive Fuzzy Clustering

This subsection presents the simulation results for AFC and LS. In these simulations, we first show the performance of AFC when there are no colluding measurements, because by using a single centralized device it is easy to prevent the nodes from having multiple colluding measurements. We also show results when there are multiple colluding measurements to illustrate the robustness of AFC under certain attacks. The robustness of AFC makes it also attractive to other ranging based localization schemes which are easier to be attacked by malicious nodes.

The settings of the simulations are as follows: the single node with an unknown location is located at the center of a 400m×400m target field; for each ranging measurement, the centralized device is randomly placed in the target field; every ranging measurement contains either a large error or a small error; the large ranging errors are uniformly distributed among [1m, max], in which max is the maximum value for that uniform distribution; for small ranging errors, the error is randomly picked from our long distance ranging experiments shown in Section .A. We do not consider severely underestimated errors in this set of simulations, because no severely underestimated error is observed during all the experiments in our experimental Thunder system tests discussed in Section .A. All the points in our following figures are computed from 1000 trials. δ ranging measurements are obtained, and the ratio of large ranging errors is σ.

We first show in Figure 6 the performance of LS and AFC when there are no colluding measurements. Figure 6(a) illustrates the impact of σ when δ=20. Almost all the average localization errors and standard deviations of AFC are below 20cm, when $\sigma \leq 60\%$. The performance starts to degrade quickly when σ □ 70%, because the number of ranging measurements with small errors becomes too small. In reality, unless the environment is extremely complicated with too many obstacles, it is not common to have σ□70%. For example, in our long distance ranging experiments shown in Section .A, σ is below 5%. Figure 6(b) illustrates the impact of different max values. The average error of AFC is very stable as the max value increases. It starts to have some sporadic large localization errors only when max is greater than 80 meters. Figure 6(c) shows the performance of LS and AFC with different values of δ. AFC starts to behave well when δ= 9. After that, almost all the average errors and standard deviations of AFC are below 20cm. Figure 6(d) shows the resistance level of AFC with different values of δ. Here we define the resistance level as the maximum value of σ to keep the average localization errors below 2.5m. The resistance level of AFC increases as δ increases. AFC can provide a resistance level of 60% when δ=15. The resistance level goes up to 81% when δ=100. So increasing the number of ranging measurements is also an effective way to further improve robustness of AFC.

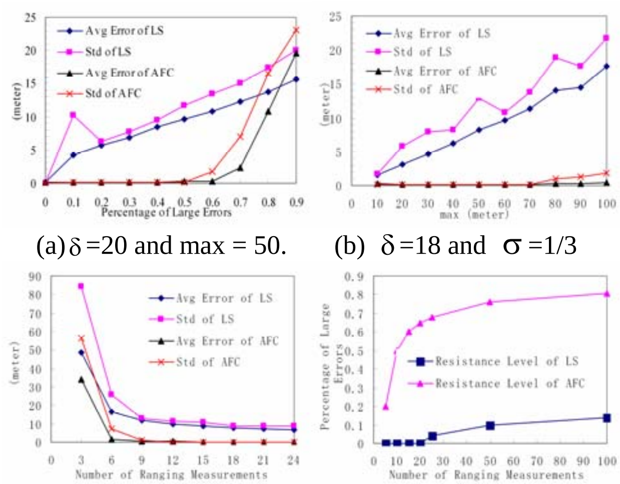

(a) δ =20 and max = 50. (b) δ =18 and σ =1/3

(c) σ =1/3 and max = 50. (d) max = 50

Figure 6: Performance of LS and AFC when there are no colluding measurements.

Figure 7 shows the performance of AFC when there are multiple colluding measurements. ε is the percentage of colluding measurements. These colluding measurements tend to push the localization to the same wrong location which is *d* meters away from the real location. Figure 7(a) shows the impact of ε. The performance of AFC degrades only when ε > 40%. Otherwise, the average localization errors are smaller than 20cm. This performance is really attractive, because no algorithm can find the correct location with high probability if ε ≥ 50% unless extra information can be provided. Figure 7(b) shows that AFC is also quite stable when the value of *d* changes. The average errors are under 20cm almost all the time.

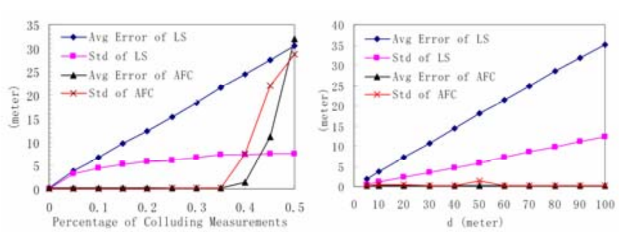

(a) *d* =50 and δ =20 (b) δ =20 and ε =35%

Figure 7: Performance of LS and AFC when there are colluding measurements.

From the simulation results, it is clear that the performance of AFC is highly satisfactory. Although we do not remove any incorrect possible locations in Phase 1, AFC can always obtain very accurate location with the average errors under 20cm when σ ≤60% if it obtains 20 ranging measurements. Also, AFC is very resilient to colluding measurements. It supports highly robust localization when the percentage of the colluding measurements is below 40%. Moreover, it is very efficient as described in Section .A. With this efficient and robust localization algorithm, Thunder can be used in complicated environments with various obstacles, which may cause a high percentage of large ranging errors.

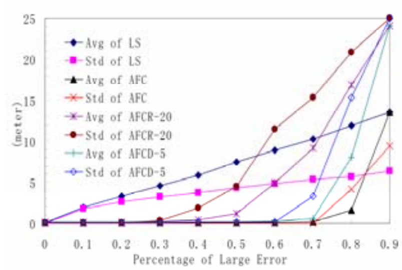

Figure 8: Performance of LS, AFC and its varieties when δ = 100 and max = 50.

One main drawback of AFC is that the computation overhead and the memory requirement increase fast as δ becomes larger. For example, ifδ =100, to store about $2\times C_{100}^{2}$ possible locations, it requires the memory of 80Kb, if each possible location takes 8 bytes. Here, we provide two possible changes to AFC to reduce the computation overhead and the memory requirements when δ becomes large. The first one is to randomly choose φ ranging measurements and then use AFC to compute the location by only using these φ ranging measurements. The second one is to first divide all δ ranging measurements into ρ subsets equally. Then we use AFC to compute a location for each subset. The location with the maximum density among all these computed locations is selected as the final location. We use AFCR-φ to denote the first method and we use AFCD-ρ to denote the second method.

Figure 8 shows the performance of LS, AFC and its varieties when δ= 100. AFCR-20 has the least computational overhead and memory requirements, however it only works well when σ□50%. AFCD-5 has about 4 times more computational overhead and similar memory requirements compared to that of AFCR-20, while it provides robust localization even when σ=70%. AFC requires about 25 times the computational overhead and memory requirements as AFCR-20 does. However, it is still resilient when σ=80%. The average error of AFC is about 1.5m when σ=80%. This is reasonable considering that only 20 out of 100 measurements contain small errors and the other 80 measurements contain large errors.

. Experimental System Implemen- tation

To verify the feasibility of Thunder to support long distance acoustic ranging, we developed an experimental Thunder system on the Mica2 platform. Due to the lack of a powerful enough radio transmitter, our experimental Thunder system is slightly different from the previously described Thunder system.

In our experimental Thunder system, we use a Mica2 as the radio transmitter. Because of its limited radio range, we need to flood the radio signal to all the sensors in the field to make them prepare for the incoming acoustic signal. So, we can not directly use the time difference of arrival between the radio signal and the acoustic signal due to message delays in the process of flooding. In our implementation, we use the time difference between the time when the acoustic signal is broadcast and the time when the sensor receives the acoustic signal to compute its distance to the speaker. Therefore, time synchronization is necessary in our experimental Thunder system. We use the time synchronization module [13] developed by Vanderbilt University. Formula (7) shows how to calculate the distance:

$$Dis = (t_{receive} - t_{broadcast}) \times Speed_{sound} \qquad (7)$$

In Formula (7), *Dis* denotes the distance from the sensor to the speaker, $t_{receive}$ denotes the time when the sensor in the field receives the acoustic signal, $t_{broadcast}$ denotes the time when the sound is broadcast from the speaker, and $Speed_{sound}$ denotes the speed of sound.

In order to get $t_{broadcast}$, we put a Mica2, which we call BMica2, quite close to the speaker to mark its location as that of the speaker. We use the time when the BMica2 detects the acoustic signal as $t_{broadcast}$. After each sound blast, the Mica2 radio transmitter floods $t_{broadcast}$ and the location of the speaker to the whole field. Based on this information, the motes in the field can compute their distances to the speaker.

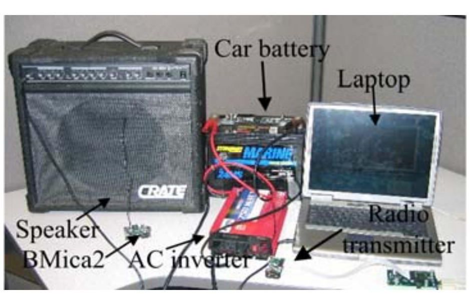

Figure 9: Our centralized device

Figure 9 shows the centralized device we use for our experimental Thunder system. The car battery is used to generate 12V DC. The AC Inverter is used to invert 12V DC to 120V, 60Hz AC, which serves as the power source for the speaker and the laptop. The laptop is used as the command center to send commands to the radio transmitter through the programming board, to generate the sound signal to the speaker, and to gather experimental data.

.A. Acoustic Signal Detection

We use the hardware phase-locked loop tone detector on the Mica sensor board to detect acoustic signals. The output of the tone detector is either 0 or 1, in which 0 means that an acoustic signal within its effective frequency range is detected and 1 means that no acoustic signal within that range is detected.

From our experiments, we find that the effective frequency range for most of the Mica sensor boards is between 4.3kHz and 5.1kHz, which is slightly different from the specification. In our implementation, we use a 4.7kHz acoustic signal which lasts for 100ms. To avoid random false detection of the tone detector caused by background noise, we accumulate the sampling results within a certain window size to see whether the number of 0s exceeds a certain threshold. If it does, the sensor records the time when it exceeds as the detection time of the acoustic signal.

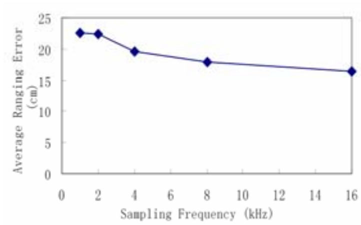

Figure 10: Average ranging errors with different sampling rates

We find that a 4kHz sampling rate is high enough to achieve sub-meter accuracy and a higher sampling rate does not help much in improving accuracy for long distance ranging. Figure 10 shows the average ranging errors with different sampling rates of the tone detector from our experiments. (We use 5 sampling rates, 1kHz, 2kHz, 4kHz, 8kHz, and 16kHz.) In these experiments, we put a Mica2 mote 2.44 meters away from the speaker and use Formula (7) to compute its distance. We have conducted the experiments 25 times for each sampling rate and plot the average errors. As shown in Figure 10, as the sampling rate increases, the average error decreases. However, the difference is very small. The average ranging error of 4kHz is only about 3cm worse than that of 16kHz, which is negligible in long distance ranging. Moreover, if the sampling rate is too high, it places negative impact on other modules. For example, the radio module does not work properly if we set the sampling rate to 16kHz. We used a 4kHz sampling rate for all the following experiments.

.B. Gain Value Adjustment

One big challenge in long distance acoustic ranging is the hardware saturation of the tone detector on the Mica sensor board caused by a strong sound signal which we use for ranging. This problem has not been addressed by previous work to the best of our knowledge. If a tone detector is saturated, it is unresponsive to acoustic signals. Until the tone detector hardware is improved, each time we do acoustic ranging, we need to first adjust the gain values of tone detectors to avoid saturation. Based on our extensive experiments, we developed a technique called the Three Phase Adjustment (TPA), which adjusts the gain value dynamically and solves this problem efficiently.

In TPA, each tone detector can only choose among three possible gain values, v1, v2 and v3. v1 is the least sensitive gain value that permits using as powerful sound source as possible. v3 is the most sensitive value. v3 should be resilient enough to environmental noise and also should be as sensitive as possible to support longer distance ranging. v2 is used to address the unreliability issues of tone detectors and its value is between v1 and v3. In Mica2 sensor boards, gain values of the tone detector range from 1 to 125. 1 means the least sensitivity and 125 means the most sensitivity. So we choose 1 for v1. And we choose 70 rather than 125 for v3, because the gain value of 125 is too sensitive, and it is vulnerable to environmental noise, while the gain value of 70 can effectively eliminate environmental noise and is sensitive enough for long distance ranging. Table 1 shows how resilient the gain value of 70 is for different kinds of sound. We choose 30 for v2.

Sound source	Resilient?	Sound source	Resilient?
Bird chirpings	Yes	People clapping	Yes
Wind noise	Yes	Car engine	Yes
Foot steps	Yes	Car horn	Yes
People talking	Yes	Helicopter	Yes

Table 1: Resilience of the gain value of 70

In TPA, we need to broadcast the sounds of the same strength for ranging two times to let tone detectors choose their appropriate gain values. Before the first sound blast, the gain value of each tone detector is set to v1, in which case, no tone detector is saturated by using our speaker. Then the tone detectors which detect the first sound blast set their gain values to v1 while others set their gain values to v2. If the tone detectors which set their gain value to v1 after the first sound blast also detect the second sound blast, the gain values remain v1 otherwise they are set to v2. Note that tone detectors are not reliable. It happens that a tone detector with the same settings sometimes can detect the acoustic signal while sometimes not. So if a tone detector with a certain gain value can detect the acoustic signal occasionally, it is safer to use a more sensitive gain value to make sure that the tone detector can detect the later acoustic signals. If tone detectors which set their gain values to v2 after the first sound blast detect the second sound blast, the gain value is v2 otherwise it is v3.

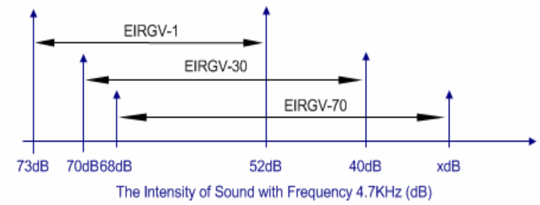

Figure 11: Effective sound intensity ranges of gain values 1, 30 and 70 for a typical tone detector

We define the Effective Sound Intensity Range of Gain Value *n* (EIRGV-*n*) as the intensity range within which a tone detector with gain value *n* can detect the sound without saturation. If the sound intensity is beyond EIRGV-*n*, the tone detector with gain value *n* is saturated, and if it is below EIRGV-*n*, the tone detector can not detect the sound. We show EIRGV-1, EIRGV-30 and EIRGV-70 for a typical Mica2 sensor board in Figure 11. In Figure 11, 73 dB is the maximum intensity of the sound with frequency 4.7KHz from our speaker. Because the sound level meter we use for our experiments can not measure the sound intensity below 40dB, we are unable to get the lower bound of EIRGV-70. Although we can see from Figure 11 that the union of EIRGV-1 and EIRGV-70 already covers the range from the maximum sound intensity of our speaker to the lower bound of EIRGV-70, it is not safe to use only these two gain values due to unreliability and variability of tone detectors and we observe some saturation when only these two values are used. In other words, if tone detectors were reliable enough, Two Phase Adjustment maybe enough, in which we only need to broadcast the sound one time before we perform acoustic ranging and each tone detector can set its gain value to 1 or 70 based on whether it detects the sound signal. The intermediate value of 30 can be used, when a tone detector is not stable with the gain value 1 and is likely to be saturated with the gain value 70. By using TPA, we do not have the saturation problem in our experiments.

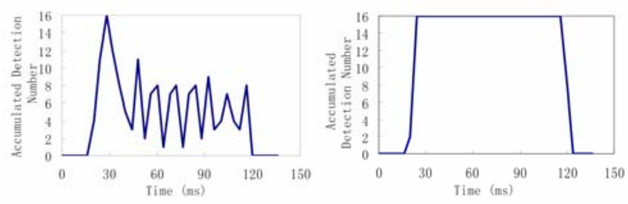

(a) (b)

Figure 12: (a) Unreliable detection of an acoustic signal. (b) Reliable detection of an acoustic signal

Another important issue in gain value adjustment is that we should set a wide enough window size and a strict enough threshold. Figure 12 shows the accumulated detection numbers of the tone detector of the unreliable detection and the reliable detection of a 4.7kHz acoustic signal over time, respectively. We accumulate 16 outputs of the tone detector into one number. The number is incremented by 1 if the output of the tone detector is 0, which means it detects the sound. As we can see from the figures, when an acoustic signal is detected reliably, it is always continuously detected. But when an acoustic signal is detected unreliably, it is detected by chance. We should avoid unreliable detection during ranging, because it results in uncertainty of the detection and may cause large errors. By setting a wide enough window size and strict enough threshold, the unreliable detection during the gain value adjustment is perceived as no detection, which makes the tone detector choose a bigger gain value as shown in TPA. By using a bigger gain value, the tone detector is more sensitive and is more likely to detect the sound with the same intensity.

.C. Angle of Sound From Speaker

We find that the angle of sound from our speaker is quite limited. As shown in Figure 13, α is relatively small. It is about 40 to 50 degrees. Normally, sensors within that angle can detect the arrival of sound signal with low ranging errors of under 50cm. But sensors outside that angle may not detect it, because the acoustic signal outside that angle is weakened dramatically, or are very much more likely to have severe late responses due to echoes. These large overestimated ranging errors can be bigger than 20 meters. We call the area within the angle α a Reliable Area and the area outside the angle α an Unreliable Area. In order to let each mote in the field have at least one correct ranging measurement to the speaker at one location, all the motes in the field need to be covered by the Reliable Area at least once for each location where the speaker broadcasts sound.

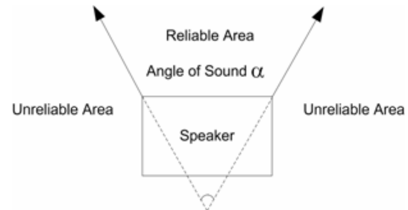

Figure 13: Angle of sound from a normal speaker.

The best solution to this problem is to use an omni-directional speaker. In our implementation, we used a narrow angle speaker, but we broadcast the sound several times at one location in different directions to satisfy the requirement that the union of the Reliable Areas can cover the whole sensor field.

. Performance Evaluation

We first perform long distance acoustic ranging experiments up to 152 meters to determine effective ranging distances. Then we perform complete localization experiments in a parking lot. The localization errors from our experiment are about 1 meter.

.A. Long Distance Acoustic Ranging

To test the feasibility of long distance acoustic ranging and the maximum ranging distance our speaker can support, we put 10 Mica2 motes in a line in front of the speaker. The experiments are done on a small lane with people walking through. The nearest mote is 15.24 meters away from the speaker, and the adjacent motes are also 15.24 meters apart. So the furthest mote is 152.40 meters away from the speaker. We conduct the experiment 17 times.

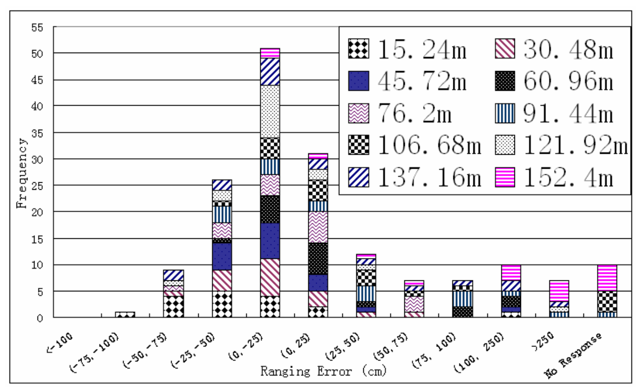

Figure 14: Distribution of ranging errors

Figure 14 shows the distribution of ranging errors from our experiments. (x, y) on the x axis means errors between x cm and y cm. No response on the x axis means that the mote fails to recognize the acoustic signal and does not get ranging estimation. Frequency on the y axis denotes the absolute number of ranging errors within (x, y). The motes with different distances to the speaker use different patterns to show the error distributions of each mote.

As we can see from the figure, the majority of the ranging errors are within 25 cm. Overall, when the mote's distance to the speaker is further away, it is more likely to have late response or fail to recognize the acoustic signal due to weakened acoustic signals. Also, no severe early detection happens during this set of experiments. In fact, we do not observe any severe early detection in all our ranging and localization experiments. A tone detector with the gain value 70 is very effective to filter out environmental noise. Even if there is some sporadic detected high frequency noise, it is filtered out by accumulating the sampling results within a certain window size.

Another interesting observation is that the majority of the experiments get underestimated distances which are a little shorter than the actual distances. The main possible reason is that all the distances are computed based on the BMica2. If the BMica2 gets a late response, which means that it uses a timestamp larger than the actual one, all the motes in the field are inclined to have shorter ranging estimations as shown in Formula (7). Another possible reason is that the speed of sound we use may not be precise. The errors caused by using an imprecise speed of sound are more obvious in long distance ranging.

Figure 15: Average ranging errors

We show the average ranging errors and the standard deviations with different distances to the speaker in Figure 15. We ignore motes with no responses in the computations. Both the motes with distances 91.4m and 137.16m to the speaker get one severely overestimated ranging error of over 7 meters, which increase their average errors to about 0.9m. The mote with distance 152.4m has the average error of nearly 2 meters, and it also fails to recognize the acoustic signal almost one third of the time. This indicates that 152.4 meters has exceeded the maximum effective ranging distance of the speaker.

.B. Medium Scale Acoustic Localization

To study the effectiveness of the complete localization solution, we deploy 18 Mica2 motes in a 24.4m × 68.6m area of a parking lot. We put these motes in 3 columns. The adjacent columns are 12.2m apart, and each column contains 6 motes. The adjacent motes in a column are 13.7m apart. The focus of this set of experiments is to validate the main idea of Thunder of using a single powerful centralized device for localization. We do not test the robustness of Thunder by using AFC in the following experiments, which is already validated through simulations in Section □. So, in the following experiments, the speaker only broadcasts sound at three corners of the field, but in order to cover the whole area, the sound is broadcast several times towards different directions at each location. In this way, a mote can get the ranging measurements to the speaker at three different locations at most. For a single trilateration at each mote, the Nelder Mead method [15], an optimization approach for nonlinear functions, can be applied to compute its location.

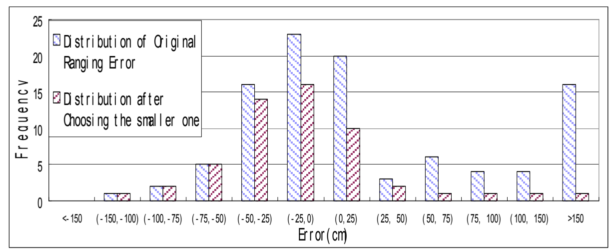

Figure 16: Distribution of ranging errors during localization experiment in the parking lot

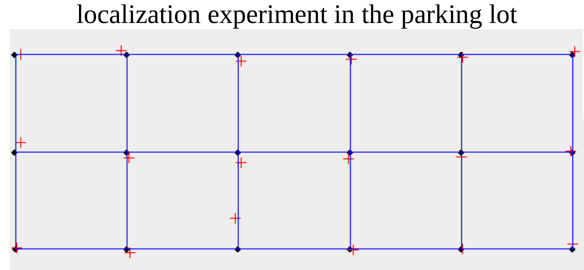
Figure 17: Localization results by broadcasting sound two times at each location.

In our first experiment, we broadcast the sounds for two times towards different directions at each location. We first show the distribution of the original ranging errors during the localization experiment in Figure 16. Overall, the majority of the original ranging errors are within 50cm. The distribution of the original errors in this figure is quite similar to that in Figure 14, except that Figure 16 contains a higher percentage of severe overestimated errors which are greater than 1.5m. That is because in our ranging experiments, all the motes are put in a line in front of our narrow angle speaker and they are in the Reliable Area of the speaker, except the furthest one which exceeds the maximum effective ranging distance of the speaker. But in the localization experiments, the motes are spread out. During each sound broadcast, the motes in the Unreliable Area are more likely to have severe late responses, causing a higher percentage of severe overestimated errors. This should not occur if we can use an omni-directional speaker.

In our implementation, if a mote obtains more than one ranging measurement to the speaker at the same location, it chooses the smaller one. By doing this, most severely overestimated errors are filtered out as shown in Figure 16. Then the Nelder Mead method [15] is applied to compute a node's location. The localization results are shown in Figure 17. The points in the figure are the actual locations and the crosses are the localization result. As shown in the figure, most motes compute their locations very close to their actual locations. The average localization error in our experiment is 1.19m. The minimum is 0.19m and the maximum is 3.96m.

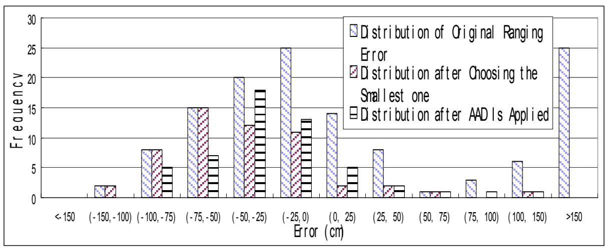

Figure 18: Distribution of ranging errors during localization experiment in the parking lot

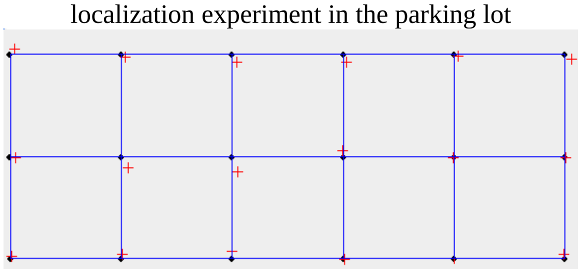

Figure 19: Localization results by broadcasting sound three times at each location.

We perform another experiment to see the effects of increasing the number of sound broadcasts at each location. This time, at each location, we broadcast the sound three times, each in a different direction. Thus, each node in the field has the chance to get more ranging measurements to the speaker at one location. The distribution of the original ranging errors shown in Figure 18 gets a higher percentage of underestimates compared to that in Figure 16, because some more late responses in the BMica2 happen in this experiment which shortens the measured distances. After each mote chooses the smallest value as the ranging measurement if it gets more than one range to the speaker at the same location, all the severe overestimated errors are filtered out as shown in Figure 18. Figure 19 shows the localization results. By broadcasting the sound three times at each location, we get results with only a little improvement. The average localization error becomes 1.10m. The minimum is 0.07m and the maximum is 3.41m.

The main reason that we do not get significantly improved results by increasing the number of sound broadcasts at each location is that although we can eliminate all the severe late responses by increasing the number of sound broadcasts, we can not eliminate the underestimated measurements by always choosing the smallest value. The best solution is to have a BMica2 with a more reliable tone detector. But due to our hardware limitations, we use a simple algorithm called Averaging After Discarding (AAD) to alleviate this problem. In AAD, if a mote has more than one range to the speaker at the same location, we discard the ranging results which are greater than the smallest value by a certain threshold. Then we use the average of the remaining measurements as its final ranging measurement. In this way, we can filter out the severely overestimated ranges and also reduce the possibility of having underestimated ones. As shown in Figure 18, by using AAD we obtain fewer underestimates. The average localization error of broadcasting three times at each location becomes 0.96m. The minimum is 0.10m and the maximum is 2.68m.

Average Time	Minimum Time	Maximum Time	Standard deviation
87ms	66ms	646ms	78ms

Table 2: Computing time of using the Nelder Mead method on Mica2 motes

One important issue of using the Nelder Mead method to compute the locations is that it is not guaranteed to find the global optimal point from some starting points. However, we can try other different starting points if it fails to find the global optimal point from some starting points, and finally find the global optimal point. Table 2 shows the computing time of using the Nelder Mead method on Mica2 motes in our previous experiments. During these experiments, over 95% of the motes succeed to find the global optimal locations from the first starting point.

.C. Summary of the Experiments

Both the long distance ranging experiments and the medium scale localization experiments validate the main idea of Thunder of using a single powerful centralized device for localization. The ranging experiments show the promising ranging distance that a centralized device can support. A speaker whose maximum sound intensity is 73dB, which is the same magnitude level as street noise in a city, is used for ranging. It can reliably support ranging up to 137 meters. We envision that with a more powerful speaker, it is possible to do acoustic ranging up to several hundred meters. Our localization experiments show the high efficiency of using a powerful centralized device for localization. For many systems we only need to move the centralized device to three different locations to localize a medium scale WSN and the average localization errors are only about 1 meter. When the sizes of WSNs grow larger and the environment becomes more complicated, ETS and AFC can be applied to provide scalability and robustness correspondingly.

Because other papers have already discussed the sources of errors in acoustic ranging, such as [9], here we do not go into the details of the error sources. We simply mention the following 5 possible error sources in our experiments: time synchronization, unreliable tone detector, multi-path signals, sampling frequency and the imprecise speed of sound. Among these, multi-path signals cause most of the large errors in our experiments.

. Related Work

In this section, we first summarize previous acoustic localization schemes proposed so far. Then we briefly discuss other localization approaches.

.A. Acoustic Localiation Schemes

Several acoustic localization schemes [9][21][23] are proposed due to the high accuracy of acoustic ranging [5][22]. The Cricket location support system [21], which is designed for context aware indoors deployment, can achieve errors of tens of centimeters by using ultrasound transceivers. It uses pre-installed anchors which know their own locations, and other static or mobile nodes in the building do acoustic ranging to these pre-installed anchors to localize themselves. AHLoS, proposed by Savvides et al. [23], uses Medusa nodes with multiple ultrasound transceivers to achieve localization errors within 20cm using a certain percentage of anchors. It uses peer-to-peer ranging and has the great potential to be used for outdoor ad hoc WSNs. However it is limited by short ranging distance for each pair of Medusa nodes. Even for the second generation Medusa nodes, the ranging distance is only 10-15 meters.

From our point of view, the acoustic localization scheme described in [9] is quite similar to AHLoS and it is designed for outdoor ad hoc WSNs. The main difference is that [9] uses audible sound instead of ultrasound. One main advantage of using audible sound is that it is much easier to generate omni-directional sound and typically has longer ranges. However, the audible sound (if too long) may annoy people (in our case the sound only lasts 100ms), and may not be suitable for hostile environments. [9] also uses peer-to-peer ranging with an extra buzzer unit and a 9V battery for each node and requires a certain percentage of anchors. From their medium scale experiments, errors of 2.5m are reported by using a non-scalable centralized Least Square Scaling algorithm. But, when the decentralized Least Square Scaling algorithm is used, errors go up to 9.5m, mainly due to the aggregation and amplification of errors during the iterative process.

.B. Other Localization Schemes

The global positioning system (GPS) [18] uses Time of Arrival (TOA) techniques to infer its location based on the speed of the radio signal. Although GPS is in wide use in both military and civil applications, it is not desirable to put it on every sensor due to its high cost. Also, the energy consuming electronics in GPS limit its use in energy constrained devices.

Ranging techniques based on Received Signal Strength Indicator (RSSI) have been extensively studied [1][12][19][24]. To avoid using anchor nodes, several localization approaches [19][24] based on RSSI use a mobile beacon. Using these approaches, errors of several meters are reported. Although some techniques such as parameter calibration and two-phase refinement positioning can be used to reduce the errors, these approaches are not used in practice due to their low accuracy from irregularity of signal propagation, multi-path fading and background interference in the real world. In [4], Elnahrawy et al. present the inherent limitations of using RSSI.

Besides these range-based localization approaches, many range-free localization approaches are also proposed, mainly to eliminate expensive extra hardware. In [2], Bulusu et al. propose a localization scheme called Centroid, in which each node computes its own location by computing the centroid of its connected anchors. Localization errors of 1.8m are reported in a 10m $\times$ 10m square with four anchors. In [6], He et al. propose an area-based range-free localization scheme, called APIT. Each node decides its location by learning whether or not it is inside a triangle of multiple combinations of the anchors. Another approach in range free localization is DV based localization [17]. It mainly uses multilateration to compute the nodes' locations by using the hop count from anchors to the nodes and the hop distance estimates. Errors from 20% to 150% of radio range are reported from their simulations. Many of these range-free localization schemes, especially DV based localization and its variants, require many peer-to-peer messages and consume significant energy.

Spotlight [25] is a newly emerged technology that uses a centralized device to scan the sensor field. It provides both high-accuracy and low-cost localization for WSNs. But, there are still issues to be solved in some application scenarios, such as providing robust localization for environments that have plants or other obstacles that block the laser beams used for Spotlight.

. Conclusions and Future Work

In this paper, we have presented a practical acoustic localization scheme called Thunder for outdoor ad hoc WSNs. Although the basic idea behind Thunder is simple, it exhibits many nice properties by stripping most of the complexities and hardware requirements from each node to a single powerful centralized device. With this asymmetric architecture, no extra hardware is required to current popular senor boards and also no in-field anchors are needed. It is fast to localize, easy to use, requires zero cost from the motes in terms of in-network communication and supports 3D localization. It can localize a very large WSN efficiently. Further, it provides robust localization without sacrificing efficiency, which enables it to be used in complicated environments with many obstacles. We developed the experimental Thunder system on the Mica2 platform to verify the idea of Thunder. Localization errors of about 1 meter are achieved from our experimental Thunder system. One possible drawback of this scheme is the need for a loud sound sent multiple times which may inhibit its use in hostile environments. However, the sound is not as loud as the name of the scheme—Thunder. For example, the maximum intensity of the sound used in our experiments is 73dB, but it supports effective ranging up to 137 meters. Further, the sound only lasts 100ms.

As future work, we will study the impact of wind on long distance acoustic ranging and develop approaches to reduce the impact.

. Acknowledgements

This work was supported by the NSF grants CCR-0329609, CCR-0325197, and CNS-0435060. Additionally, we would like to thank Gang Zhou, Pascal Vicaire, Tian He, Radu Stoleru, and Lin Gu for providing valuable comments and suggestions. We also would like to thank Second Face Band at UVa for lending us the speaker.

References

[1] P. Bahl, V.N. Padmanabhan, “RADAR: An In-Building RF-based User Location and Tracking System”, in Proceedings of InfoCom, 2000.

[2] N. Bulusu, J. Heidemann, D. Estrin, “GPS-less Low Cost Outdoor Localization for Very Small Devices”, in IEEE Personal Communications Magazine, Special Issue on Smart Spaces and Environments, 2000.

[3] O. Cramer, “The Variation of the Specific Heat Ratio and the Speed of Sound in Air with Temperature, Pressure, Humidity, and CO_2 Concetration”, Journal of the Acoustical Society of America, vol. 93, no. 5, pp. 2510–16, May 1993.

[4] E. Elnahrawy, X. Li, R. Martin, “The Limits of Localization using RSSI”, in Proceedings of SECON, 2004.

[5] L. Girod, D. Estrin, “Robust Range Estimation Using Acoustic and Multimodal Sensing”, in Proceedings of IROS, 2001.

[6] T. He, C. Huang, B. Blum, J. A. Stankovic, T. Abdelzaher, “Range-Free Localization Schemes for Large Scale Sensor Networks”, in Proceedings of Mobicom, 2003.

[7] T. He, S. Krishnamurthy, J.A. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, B. Krogh “An Energy-Efficient Surveillance System Using Wireless Sensor Networks”, in Proceedings of Mobisys, 2004.

[8] R. Kershner, “The Number of Circles Covering a Set”, American Journal of Mathematics, v.61, pp. 665-671, 1939.

[9] Y. Kwon, K. Mechitov, S. Sundresh, W. Kim, G. Agha, “Resilient Localization for Sensor Networks in Outdoor Environments”, in Proceedings of ICDCS, 2005.

[10] Z. Li, W. Trappe, Y. Zhang, B. Nath, “Robust Statistical Methods for Securing Wireless Localization in Sensor Networks”, in Proceedings of IPSN, 2005.

[11] D. Liu, P. Ning, W. Du, “Attack-Resistant Location Estimation in Sensor Networks”, in Proceedings of IPSN, 2005.

[12] K. Lorincz, M. Welsh, “MoteTrack: A Robust, Decentralized Approach to RF-Based Location Tracking”, in Proceedings of LoCA, 2005.

[13] M. Maroti, B. Kusy, G. Simon, A. Ledeczi, “The Flooding Time Synchronization Protocol”, in Proceedings of Sensys, 2004.

[14] D. Moore, J. Leonard, D. Rus, S. Teller, “Robust Distributed Network Localization with Noisy Range Measurements”, in Proceedings of Sensys, 2004.

[15] J.A. Nelder, R. Mead, “A Simplex Method for Function Minimization”, Comput. J.7, 308-313, 1965.

[16] D. Niculescu, B. Nath, “Ad Hoc Positioning System (APS) using AoA”, In Proceedings of IEEE InfoCom 2003, San Francisco, CA.

[17] D. Niculescu, B. Nath, “DV-based Positioning in Adhoc Networks”, in Telecommunication Systems, vol. 22, 2003.

[18] B.W. Parkinson, J. Spilker, “Global Positioning System: theory and applications”, in Progress in Aeronautics and Astronautics, vol. 163, 1996.

[19] P.N. Pathirana, N. Bulusu, A. Savkin, S. Jha, “Node Localization Using Mobile Robots in Delay-Tolerant Sensor Networks”, in IEEE Transactions on Mobile Computing, 2004.

[20] N. Priyantha, H. Balakrishnan, E. Demaine, S. Teller, “Mobile-Assisted Topology Generation for Auto-Localization in Sensor Networks”, in Proceedings of InfoCom, 2005.

[21] N. Priyantha, A. Chakaborty, H. Balakrishnan, “The Cricket Location-support System”, in Proceedings of MobiCom, 2000.

[22] J. Sallai, G. Balogh, M. Maroti, A. Ledeczi, B. Kusy, “Acoustic Ranging in Resource-Constrained Sensor Networks”, Technical Report, ISIS-04-504, February 25, 2004.

[23] A. Savvides, C.C. Han, M.B. Srivastava, “Dynamic

fine-grained localization in ad-hoc wireless sensor networks”, in Proceedings of MobiCom, 2001.
[24] M. Sichitiu, V. Ramadurai, “Localization of Wireless Sensor Networks with a Mobile Beacon”, in Proceedings of MASS, 2004.
[25] R. Stoleru, T. He, J.A. Stankovic, D. Luebke, “High-Accuracy, Low-Cost Localization System for Wireless Sensor Networks”, in Sensys, 2006.
[26] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, D. Culler, “An Analysis of a Large Scale Habitat Monitoring Application”, in Proceedings of Sensys, 2004.
[27] K. Whitehouse, A. Woo, C. Karlof, F. Jiang, D. Culler, “The Effects of Ranging Noise on Multi-hop Localization: An Empirical Study”, in Proceedings of IPSN, 2005.
[28] R. Williams, "Circle Coverings", The Geometrical Foundation of Natural Structure: A Source Book of Design, New York: Dover, pp. 51-52, 1979.
[29] http://ceti.cse.ohio-state.edu/exscal/