

Robust Range Estimation Using Acoustic and Multimodal Sensing

Lewis Girod and Deborah Estrin

UCLA

Email: {girod,destrin}@lecs.cs.ucla.edu

Abstract

Many applications of robotics and embedded sensor technology can benefit from fine-grained localization. Fine-grained localization can simplify multi-robot collaboration, enable energy efficient multi-hop routing for low-power radio networks, and enable automatic calibration of distributed sensing systems. In this work we focus on range estimation, a critical prerequisite for fine-grained localization. While many mechanisms for range estimation exist, any individual mode of sensing can be blocked or confused by the environment. We present and analyze an acoustic ranging system that performs well in the presence of many types of interference, but can return incorrect measurements in non-line-of-sight conditions. We then suggest how evidence from an orthogonal sensory channel might be used to detect and eliminate these measurements. This work illustrates the more general research theme of combining multiple modalities to obtain robust results.

1 Introduction

One school of thought in the robotics community contends that an essential property of a robot is its ability to operate in an environment that has not been modified to support the robot's activities. While this is a reasonable guiding principle, it is often useful to relax this assumption and allow the *robot* to modify its environment in order to help it or its collaborators accomplish their tasks.

Modification of the environment for communication and collaboration, sometimes known as *stigmergy*, is a very powerful tool used to great advantage by some of the most successful animals on earth.[1] In the context of robotics, the idea of stigmergy can be taken a step further if the robot implants active devices into the environment that can enable it to operate more efficiently. For example, compared with instrumenting the environment by installing a localization system, suppose the robots themselves could "install" a localization system by depositing small wireless devices throughout the environment.

With the decreasing cost of low-power microcontrollers and small integrated radio packages, it is becoming increasingly practical to design small wireless sensor-processors

that robots can drop or affix to strategic locations in the environment. Such a project involves many challenges, from developing the mechanisms that enable the robot to place the devices to designing low-power systems that can perform a useful task for days, weeks or months after being deployed in an ad-hoc and possibly error-prone fashion.

When developing an application intended for ad-hoc deployment, it quickly becomes clear that if the devices lack knowledge of their position, it significantly decreases their utility. With fine-grained localization information, they can act as a positioning system for nearby robots, or perform collaborative sensing and signalling, such as beamforming[2]. Without localization (apart from the coarse form implied by radio proximity) it is difficult to meaningfully locate small robots or combine and process sensor readings in a meaningful way.

Previous work in the area of coordinated localization systems has addressed fine-grained[3, 4], ad-hoc deployable[5, 6], distributed[3, 4, 5, 6], and decentralized[5, 6]. However, this work represents the first steps towards a system that addresses all of these characteristics at once. Our approach is characterized by three principles:

1. For every sensory system, there exists a set of environmental conditions that will confuse it, and a subset of those in which it fails to identify that it is confused.
2. Some sensory modalities are "orthogonal" to each other, meaning that their sets of failure conditions (as described above) are largely disjoint.
3. Orthogonal modalities can identify each others' failure modes, and thus improve the data quality through coordination and communication, with significantly less effort relative to the effort required to incrementally improve the sensors on their own.

For example, consider a system composed of many standalone ranging units and a few ranging units with cameras. In general, acoustic ranging performance suffers when the "line of sight" (LOS) path is obstructed. Acoustic range measurements in obstructed conditions often consistently detect longer reflected paths, leading to unbounded range error. Because they measure the long path consistently, it can

be very difficult to identify these errors based exclusively on analysis of the acoustic data.

However, suppose that in our example each camera’s field of view contains several ranging units, which might be identified by a characteristic pattern strobed on an IR LED. Any ranging unit that the camera can see has a high probability of LOS to the camera, and thus in those cases an accurate range can be determined with acoustics. Additionally, using angular displacement, a camera can estimate the range between any two ranging units in its field of view. By using the relatively coarse angular information from the camera, ranging units would be able to identify and ignore large errors resulting from obstructed conditions.

This approach can be seen as a form of “sensor fusion”. (See Roumeliotis[7] for a good example of sensor fusion.) What is novel in this case is that the sensors are distributed across a network, which introduces significant complexity and additional concerns: e.g. power and bandwidth conservation, resilience to temporary and permanent loss of connectivity, graceful degradation as nodes fail, and scaling issues to name a few.

In this paper we discuss an implementation of an acoustic ranging system in support of a fine-grained, ad-hoc, decentralized system. In Section 2 we discuss ranging alternatives and our reasons for selecting acoustics. In Section 3, we describe the implementation and calibration of our acoustic ranging system, and show how its properties are appropriate to our goals. Then, in Section 4 we characterize the sensor in more depth, specifically focusing on cases in which it is difficult for the sensor to identify that it is failing, and suggest some approaches for using orthogonal modalities to solve the problem.

2 Why Acoustics?

When we examine high-accuracy omni-directional ranging techniques, the most successful techniques are based on measuring the time of flight of signals. Other common techniques, such as RSSI or RF proximity estimation, tend to be highly susceptible to environmental interference and are also non-linear. While they may be useful as a supplemental technique, they are less useful as a standalone measurement. In time-of-flight approaches there are two competing technologies: those based on RF and those based on acoustics. In comparing the two alternatives, we identified three important advantages of acoustics over RF.

First, the low frequencies and slow speed of acoustic signals reduces the cost and complexity of the ranging hardware. Sampling rates of 40 or 100 kHz are sufficient to adequately recover audible acoustic signals. Further, because the sound waves travel slowly ($\approx 344\text{m/s}$), the sampling rate of a standard sound card (48kHz) is sufficient for a position sensing granularity under 1 cm. Besides the reduced cost and complexity, there may also be an advantage in terms of power requirements. Relative to RF-based approaches,

acoustic ranging shifts the power costs from the receiver to the transmitter. An acoustic receiver doesn’t need to maintain accurate clocks because it can rely on post-facto synchronization (see [8] for further analysis of post-facto synchronization), and an acoustic transducer can be operated at lower power because much of the energy to operate it can be drawn from the signal itself. These two factors might enable the receiver circuitry to power down when it is not in use with almost no loss of functionality.

Second, synchronization of the sender and receiver can easily be accomplished with radio signals. This is an important contrast to ranging based on electromagnetic signals, which must rely on a highly accurate reflection from the target in order to measure the time of flight.¹ Not only does this dramatically simplify the system, it also results in a significant reduction in error. The reason is that a large part of the error in a ranging system enters during signal detection, when noise and environmental factors conspire to muddy the signal. A system in which the target is synchronized requires only one detection phase, compared with two in the case of reflective or retransmission ranging.

Third, audible acoustic waves encompass an extremely wide range of wavelengths despite their relatively small range of frequencies. Since environmental features of a given size tend to scatter waves of the same length, this means that a wideband acoustic signal can be resilient to a wider range of scattering features. To illustrate this, if the ranging signal uses a band from 300Hz through 20kHz, the component wavelengths will range from 1m to 1.5cm. In contrast, an RF UltraWideBand (UWB) ranging system that uses the range from 2GHz to 3GHz only includes wavelengths from 30cm-10cm.

Although they have many advantages, acoustic systems do have two drawbacks. First, even over short ranges they suffer a significant local dependence on atmospheric conditions, necessitating compensation in order to achieve high accuracy. Second, they exhibit lower penetration of solid objects than RF, which can make inadvertent detection of reflected paths more common. However, despite these flaws the advantages of using acoustics made it the clear choice in this case.

3 An Acoustic Ranging System

Similar to the design of other acoustic systems such as the Active Bat[3], our ranging system operates by measuring the time of flight of acoustic waves through air. The system is composed of a transmitter and a receiver, each of which operate autonomously. The transmitter produces a

¹It is interesting to note that in the case of GPS, the receiver need not explicitly reflect a signal because it receives signals from multiple satellites, which are themselves well-synchronized. This does not effect the argument however, because it still requires multiple detections, each contributing its own error.

distinct sound (which we call a “chirp”²), and the receiver detects that chirp with a matching sensor. Concurrently with sending the chirp, the transmitter communicates a synchronization message by radio that tells the receiver what time the chirp was emitted. By comparing the time of emission with the time at which the chirp is detected at the receiver, the time of flight of the acoustic signal can be computed. Then, by using a known model for the speed of sound in air, the distance traveled by the waves can be estimated.[9]

Much of the improvement demonstrated in this work comes from the use of pseudonoise sequences as the acoustic signal. Using a matched filter, these signals can be detected with a high degree of temporal accuracy. This technique also provides substantial processing gain, allowing highly attenuated line-of-sight signals to be detected in obstructed conditions, even when a much stronger reflected component is also received. These techniques are often classified as “spread spectrum” mechanisms, because the transmission band is much wider than the bandwidth required to transmit the data.[10]

In the remainder of this section we will describe our implementation in some detail, highlighting the areas where the design is critical to reducing noise.

3.1 Hardware

Our current implementation is built from COTS hardware with a few minor customizations. It is based on a PC platform running Linux, and uses a standard sound card for emitting and capturing acoustic signals.

The emitters and sensors are readily available parts. For microphones we used cheap microphones intended for use with PC sound cards. The speakers we used are small, full-range speakers intended for use in laptops. They are manufactured by Kingstate, model KDS-27008. They are well suited to this purpose because they are small and have a very wide frequency response, from 400Hz through 20 kHz. Their rated power is 0.1W.

The most difficult component of an implementation on PC hardware is achieving synchronization between the sender and receiver. Most standard means of communication between PC’s suffer from variable latency. For example, in the case of a network packet, there may be variable latencies at the MAC layer, in the network card, in the network driver, and in context switching to and from the application processes. All of these variable latencies result in synchronization error; 30 microseconds of latency results in a 1 cm offset.

To solve this problem, we wrote a Linux device driver that captures an interrupt from a pin on the parallel port, and can immediately record a timestamp. For the experiments described in this paper, we used a wiring harness

²The term “chirp” is usually used to mean a signal with a linear change in frequency over time. In this case we use the term to mean only that it is distinctive.

connected to the parallel port (described in more detail in [8]) to achieve synchronization without exposing ourselves to packet loss in the RF channel. We are currently working to integrate this system with a small microcontroller-based radio card developed by Pister et. al.[11, 12] which provides an interrupt on packet arrival through the same parallel port pin.

Another source for variable latency is in the sound card. Commercial sound cards are not intended for microsecond accuracy. We found that although our cards exhibited relatively accurate timing when playing sounds, when recording they introduced a variable latency of as much as 500 microseconds. In order to correct this problem, we inject a synchronization pulse directly into the recorded time series, by feeding a pin on the parallel port into the microphone input.

These techniques, although crude, reduce synchronization error to approximately a 5 sample range. By taking many trials at each measurement point and performing a statistical analysis, the accuracy can be improved to better than a 2 sample range, or about 1 cm. As part of future work, a new implementation with a higher input sample rate and designed-in synchronization might improve system performance while requiring fewer trials.

3.2 Chirp Detection Algorithms

In a previous paper on this subject[13] we discussed our experimentation with different coding and modulation strategies for producing the chirp. We concluded that using a “Direct Sequence” style binary phase shift keying (BPSK) waveform worked most consistently and was the simplest to implement.

In our new design we have selected a 511 bit code from a family of binary maximal-length sequences (m-sequences). These codes have been shown[14] to have very good auto-correlation properties, and bounded cross-correlation within the family of codes. The transmitter forms the chirp by modulating the code using BPSK at a 12 kHz chip rate.

The detector uses a matched filter to determine the position of the chirp within the observed signal. Because the receiver knows the code sequence used by the sender, it can compute a reference signal that is identical to the one transmitted. This is used to compute a correlation function that describes, for each possible offset, the degree to which the observed signal at that offset agrees with the reference signal. The resulting correlation function will have a “peak” at every offset at which there was a strong correlation. The earliest peak in the function is a candidate to be the true phase offset; later peaks are likely to be the result of ringing or echoes.

The requirement that we discover the *first* peak introduces some interesting problems for the detection algorithm. The problem arises when we must distinguish the earliest “peak” from random correlations with environmen-

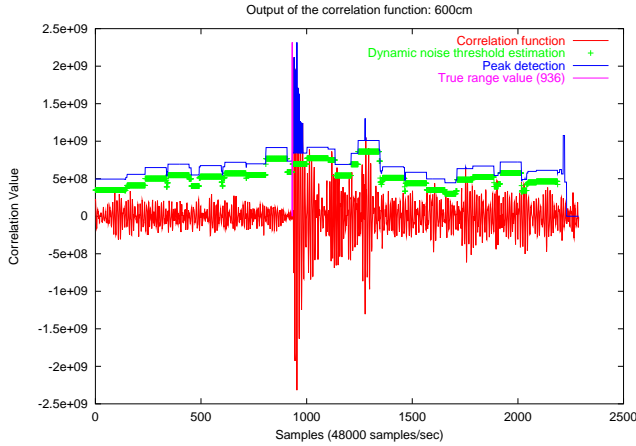


Figure 1: Output of correlation function: 600cm

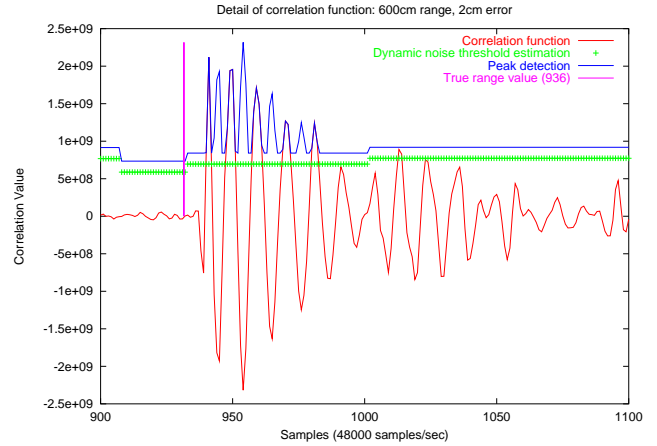


Figure 2: Detail of correlation function: 600cm

tal noise. There are various techniques for solving this problem in communications systems, although they are not generally applicable to acoustic ranging.

One common technique measures the channel at a time of known silence (e.g. before the chirp) and then uses the maximum correlation during that interval as a “noise floor”. However, this technique does not work that well in this context, because the chirp (and its flight time) is quite long and the level of environmental noise can vary significantly during that interval.

Having discovered the limitations of this approach, we have invented a somewhat nonstandard solution to this problem. Our solution computes the noise level using a second correlation operation. However, rather than correlating with the transmitter’s code, we correlate with a *different* code selected from the same family. Having similar characteristics, this code will correlate with noise in a similar way, so it will be a good indicator of the amount of noise during the whole sampling interval. However, this alternate code will not correlate well with the actual signal because of the m-sequences’ low cross-correlation properties.

Once we have performed this alternate correlation, one possible solution would be to use the maximum value of that correlation as a noise floor, and look for the first peak that rises above that floor. However, we have found through experiment that there is often a local burst of noise that will cause this algorithm to overestimate the noise floor and miss a peak. To combat this, we implemented an algorithm that computes at every point the maximum value within a specified number of samples. Peaks above that threshold function (plus an additional offset) are considered to be valid peaks.

A sample correlation function is shown in Figures 1 and 2. The x axis of the graph represents the sample index, and the y axis measures the value of the correlation at each sample offset. The dynamic noise floor estimate is shown as the disjointed line of x ’s that trace a path above the noise.

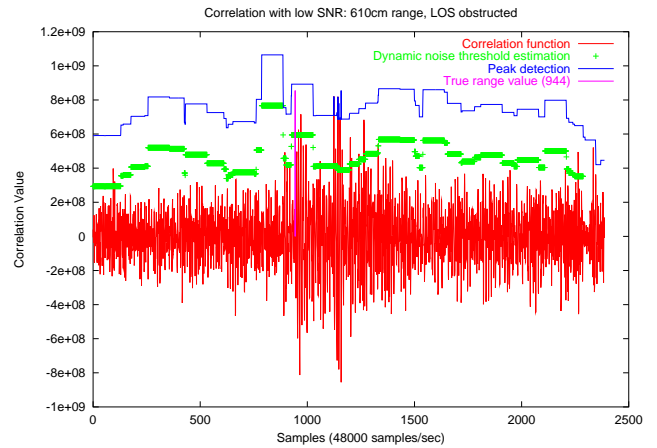


Figure 3: Correlation with low SNR

The peaks are traced by a line that follows an offset above the noise estimate. The detail (Figure 2) shows that there is approximately 2 cm of synchronization error.

While our experiments show that our detection algorithm is quite effective, it does not succeed in every case. In any given set of trials there will often be occasional trials in which the SNR is very low, increasing the probability of erroneous results. Figure 3 shows an example of such a case. By comparing the graphs, it is easy to see that the SNR is much lower than our previous examples.

In these cases, errors in peak detection can result in under or overestimation of the range. If the noise threshold is estimated too low, an erroneous early peak will be detected. Looking at Figure 3, we note that there are several instances in which noise correlates exceed the threshold prior to the “true range”: at about 100 and 900 samples. In this case it happens that our additional offset over the threshold prevented these false peaks from being detected.

Unfortunately, the same conservative estimate that

helped us prevent an underestimate in Figure 3 has instead caused an overestimate. The additional offset caused the true initial peak to be missed entirely, catching the second peak instead. As the SNR becomes lower, the probability of one of these outcomes occurring becomes higher.

In general, our algorithm is designed to favor overestimation. The reason for this choice is that regardless of the design of our detector, sufficiently obstructed conditions will always result in overestimates that cannot be detected through analysis of the signal. In contrast, since underestimates can never occur in nature, they can only be the result of a failure in the detection algorithm. Consequently, we favor overestimates in order to simplify the work of algorithms that process the range data.

3.3 Statistical Analysis

In our experiment, we performed 20 ranging trials at each measurement location. We then applied statistical analyses to each set of trials to come up with a mean value and 95% confidence intervals. In order to select an appropriate statistical algorithm, we first analyzed the potential sources of error. There are three main sources of error in our system:

1. **Synchronization error.** As we have described earlier, synchronization problems arise from hardware latency, O/S interrupt latency, and other factors. We assume that these errors can be modeled by a gaussian noise source.
2. **Intermittent detection of an attenuated peak.** If an early peak is on the verge of being detected, it may appear in some trials but not others. This can lead to a multi-modal distribution in the results.
3. **Random detection of noise correlates.** As we described in the previous section, our detection algorithm can produce incorrect results when the SNR is low. Since they are the result of correlation with noise, these values should be uncorrelated.

The first source of error can be removed by averaging.[15] However, the other two error sources result in distributions that make directly computing an average inappropriate: intermittent detection results in a multimodal distribution, and correlation with noise results in meaningless outliers.

To address these issues, we implemented a “sloppy mode” algorithm that calculates for each data point a weighted sum of all data points within a 5 sample offset. Each data point included in the sum is weighted by an estimate of the SNR for that point. Figure 4 shows four successive histograms of the sloppy mode function, for four successive measurements in a heavily obstructed case.

At 608cm range, all of the data points are concentrated in a single cluster at sample 1125, except for a single outlier at 1195. However, as the emitter moves to 610cm and

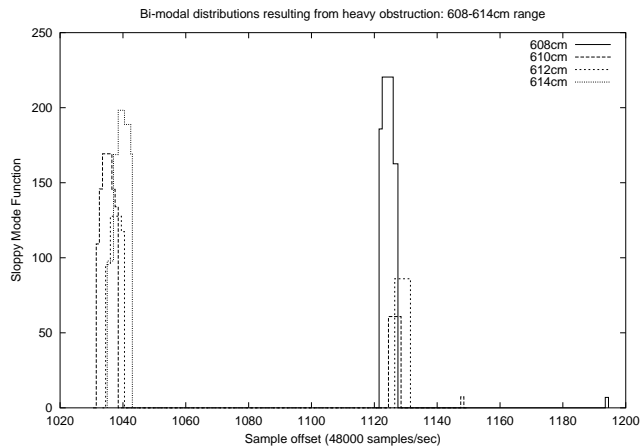


Figure 4: Bi-modal distributions

612cm, a shorter path that had previously been too heavily attenuated to be detected becomes intermittent, resulting in a bi-modal distribution. As the emitter continues to 614cm range, the new path becomes consistently detectable.

Since LOS is always the shortest path, a bi-modal distribution always indicates that LOS may be partially obstructed. Clearly averaging over the whole distribution will surely result in an incorrect answer; averaging over the cluster representing the shorter path *might* determine the correct value. But in general the knowledge that the distribution is bimodal will probably be valuable to higher level algorithms that operate on collections of range data, such as multilateration algorithms.

Our current algorithm selects the cluster containing the highest “sloppy mode” value, and then computes the mean and 95% confidence intervals on that reduced dataset. In this way, outliers are eliminated and the strongest mode is selected in the case of a bi-modal distribution.

3.4 Calibration and Compensation

For calibration of the range sensor, we take a very simple approach. If the speed of sound and the sample rate is known, then the sensor should exhibit a perfect linear relation with true distance. In this case, the only variable that must be calibrated is an additive offset; namely, the number of samples of delay after the synchronization pulse and before the acoustic chirp is transmitted. This can be easily determined by placing the emitter and receiver as close as possible and measuring the observed offset.

Once the offset of the origin has been determined, we must ensure that the speed of sound and the sample rate are correct. The sample rate is generated by the sound card, and is quite reliable. However, the speed of sound depends to a significant degree on atmospheric conditions[16]. Temperature, barometric pressure, and relative humidity must all be taken into account in order to get an accurate value for the speed of sound. Figure 5 graphically shows the effect

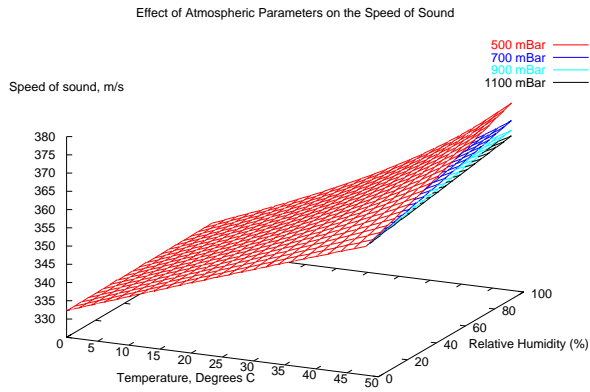


Figure 5: Effect of atmospheric parameters

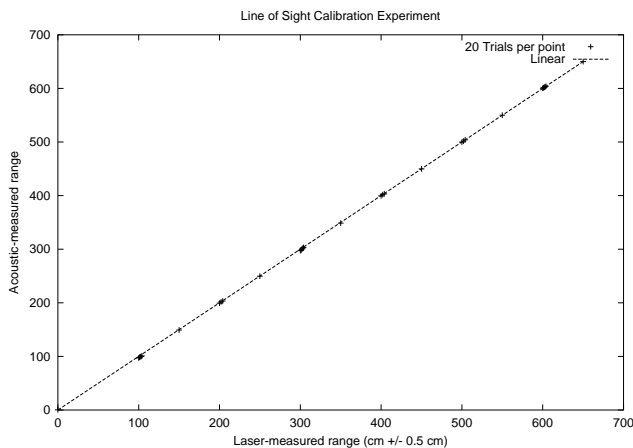


Figure 6: Line of sight (LOS) calibration experiment

of these parameters, which adds up to a fluctuation of over 10% over the full range of different values.

After establishing an origin point and the correct parameters to compute the speed of sound, we performed a calibration test. In this test, clusters of closely spaced measurements were taken at periodic intervals between 0 and 650 cm range. The results of the test are shown in Figure 6. The graph shows that the sensor behaves quite linearly. A detail shown in Figure 7 shows that under some conditions the system can be quite accurate (the errorbars represent 95% confidence intervals from 20 trials at each point.)

Figure 8 shows a different set of conditions. We believe that this unusual slope was caused by localized temperature fluctuations resulting from a ventilation duct positioned directly over the 3 meter position where these measurements were taken. In any event this condition only exists on a local scale, because globally the slope is clearly linear.

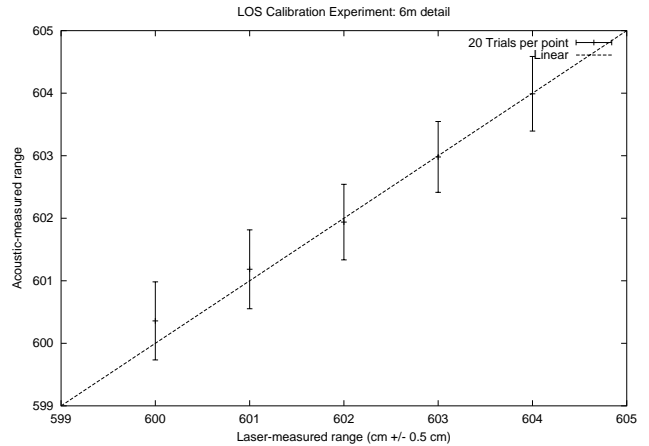


Figure 7: LOS calibration experiment: 6m detail

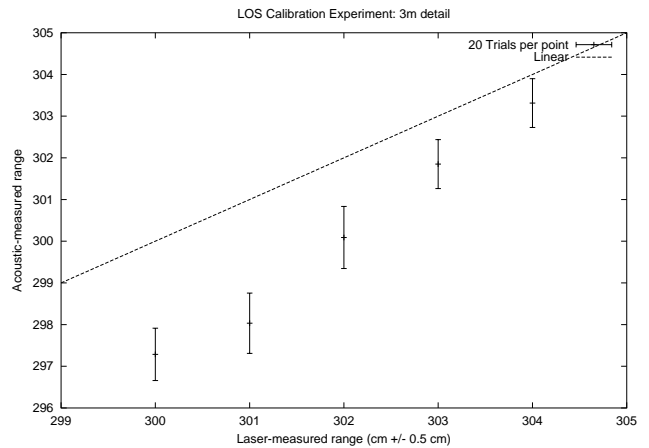


Figure 8: LOS calibration experiment: 3m detail

4 Results and Characterization

The results from the LOS calibration experiment showed a long-term linear relation, with some short-term variations that were believed to be related to local temperature fluctuation. In this section we will present a more in-depth characterization of the effect of different environmental conditions, paying particular attention to cases in which the sensor cannot estimate its error on its own.

4.1 Local Temperature Dependence

In order to test our theory that local fluctuations in temperature could result in local changes in slope, we performed an experiment in which some heated water was positioned nearby the emitter, producing a region of warm, humid air. Figure 9 shows the result of this. Shortly after the emitter enters the region of heated air, the measured range drops off the linear relation and continues at a lower slope. This corresponds well with our hypothesis, because each incremental change in position in the heated region will

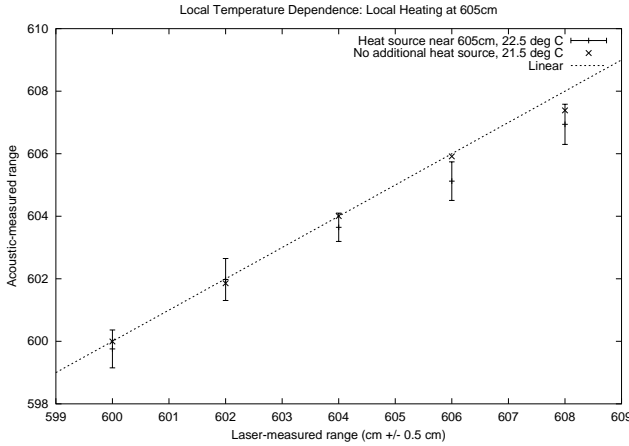


Figure 9: Local Temperature Dependence

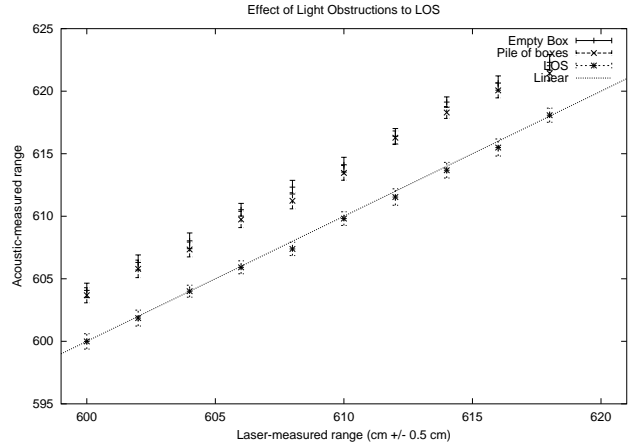


Figure 11: Effect of lightweight obstacles

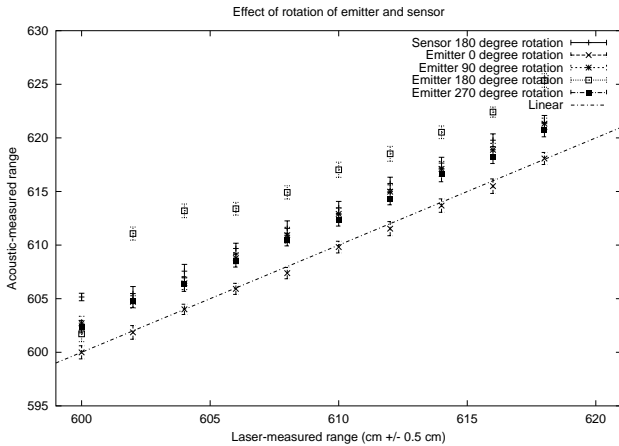


Figure 10: Orientation dependence

be undercounted in the range calculation because the sound travels slightly faster in the heated air.

There are several possible ways to address this problem. First, temperature and humidity sensors can be distributed throughout the environment to enable compensation (in fact, we have done this for our current experiments.) Second, temperature and humidity are likely to fluctuate over time, suggesting that for fixed sensor positions long-term averaging should eliminate localized fluctuations.

4.2 Orientation Dependence

In an ad-hoc deployment scenario, we assume that we cannot predict angular position with respect to other sensors. To assess our sensor's dependence on orientation, we performed some experiments in which the emitter and the sensor were rotated relative to each other.

Figure 10 shows the results of our experiment. We rotated the microphone 180 degrees, and then separately rotated the emitter 90, 180, and 270 degrees. All of the ro-

tation scenarios resulted in an offset of approximately 3cm, with the exception of the emitter rotated 180 degrees (facing away from the sensor), which exhibited a 6cm offset. In addition, the 180 degree case suffered from inconsistencies early in the trial for reasons that are not known with certainty.

In general, the offsets caused by these factors are fairly small and bounded, but they seem to be fairly consistent, which makes them very difficult to detect from analysis of the signal. In an ad-hoc setting, variations in relative orientation will be quite common, and if left uncompensated the resulting offsets may introduce significant error when supplied as the input to a multilateration algorithm.

We envision a number of possible solutions. First, because the error is bounded, it might be possible to use an optimization algorithm to determine the orientations. This solution has the drawback that it requires an expensive centralized computation. Second, the emitter could be designed to be more omni-directional. Third, a tilt sensor and a compass could be added to enable the sensor to estimate its absolute orientation. Fourth, in the event that cameras are available they might be able to visually determine orientation.

4.3 Effect of Obstructions

Obstructions to the line of sight are a critical problem for an acoustic ranging system. While in many cases these problems can be mitigated by careful placement of the system components (e.g. placing receivers on the ceiling), this is not generally an option for an ad-hoc or robot-deployed system. To assess our sensor's resilience to obstructions, we performed some experiments in which we placed various obstacles so that they obstructed the LOS path. The results of these experiments are shown in Figures 11 and 12.

In Figure 11, we did two experiments with relatively lightweight obstructions: an empty cardboard box and a stack of small cardboard boxes. Despite the fact that LOS was blocked in both of these scenarios, the average offset in

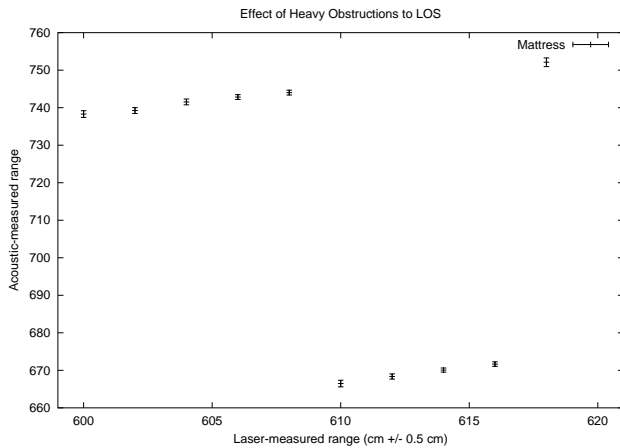


Figure 12: Effect of heavy obstacles

path length was only about 5cm. This suggests that the signal was transmitted through the obstacle, or perhaps traveled along the surface, but it did not result from a reflection.

In Figure 12, we performed another experiment with a “heavy” obstruction: a large mattress surrounded by cardboard boxes. The total size of the obstruction was on the order of 4 square meters. As is clear from the graph, the LOS path is not detectable; the receiver detects a reflected path, resulting in an offset of 140cm from the true range shown on the X axis. Another interesting observation is that there is a large discontinuity in the graph. This discontinuity is caused by changes in the multipath environment that result in the detection of different reflected paths as the emitter changes its location. For example, certain reflections might exist in one location but not in another. This results in regions of consistency (e.g. 600-607cm and 611-615cm), bounded by border regions which are likely to exhibit bi-modal distributions. In fact, the example of bi-modal data in Figure 4 was taken from this dataset.

5 Future Work

Error caused by obstructions to LOS is the most serious problem to solve in fine-grained, ad-hoc acoustic ranging. Because the magnitude of the error is unbounded, including these long paths will cause enormous problems for a multilateration algorithm. There are several approaches which might prove effective in detecting these cases.

First, detection of obstructed conditions might be possible from analysis of the signals, for example in the case that a measurement exhibits a bi-modal distribution. Signals that are the result of reflections may also exhibit higher variance, because slight perturbations in the reflecting surface can cause changes in the length of the reflected path.

Second, even if we assume that deployed ranging sensors do not move, they might be able to leverage sensors that do move (e.g. located on a robot) in order to spot cases in which

there is a sudden discontinuity. By observing a sequence of range measurements and correlating it with dead-reckoning data from the robot and a model of kinematics, it is possible in some cases to identify discontinuities might suggest that LOS is obstructed.

Third, in cases where none of the parts of the system are moving, it might be possible to have a “known good” cluster of nodes that have LOS among them. By finding ranges to each other, these nodes can establish a consistent coordinate system, and then can compare their measured ranges to external points. If when they compare notes they find that there is an inconsistency, this suggests that the path to the external point is obstructed for at least one of the nodes in the cluster. Note that the converse is not necessarily true; the appearance of consistency does not guarantee that the external point has LOS.

Fourth, by using cameras (as described in Section 1) and LEDs a node equipped with a camera might determine a set of nodes to which it has LOS. In a more complex scenario, two cameras might coordinate to use stereopsis to formulate a 3D model of the terrain and thus determine the location of obstructing features, applying the techniques of Kanade[17].

6 Acknowledgements

The authors would like to thank the Laboratory for Embedded Collaborative Systems for helpful comments and discussions regarding these ideas. This work was supported by DARPA under grant No. DABT63-99-1-0011 as part of the SCADDS project.

References

- [1] P.P. Grasse, “La reconstruction du nid et les coordinations interindividuelles chez *bellicositermes natalensis* et *cubitermes* sp. la theorie de la stigmergie: essai d’interpretation du comportement des termites constructeurs,” *Insectes Sociaux*, , no. 6, pp. 41–81, 1959.
- [2] K. Yao, R. Hudson, C. Reed, D. Chen, and F. Lorenzelli, “Blind beamforming on a randomly distributed sensor array system,” *IEEE JSAC*, vol. 16, no. 8, Oct. 1998.
- [3] A. Ward, A. Jones, and A. Hopper, “A new location technique for the active office,” *IEEE Personal Communications*, vol. 4, no. 5, Oct. 1997.
- [4] P. Bahl and V. Padmanabhan, “Radar: An in-building rf-based user location and tracking system,” in *IEEE INFOCOM*. IEEE, Mar. 2000.
- [5] N. Priyantha, A. Chakraborty, and H. Balakrishnan, “The cricket location support system,” in *Mobicom 2000*. ACM, Aug. 2000.

- [6] N. Bulusu, J. Heidemann, and D. Estrin, "Gps-less low cost outdoor localization for very small devices," Tech. Rep. CS-TR-00-729, University of Southern California, Apr. 2000.
- [7] S.I. Roumeliotis and G.A. Bekey, "An extended kalman filter for frequent local and infrequent global sensor data fusion," in *Sensor Fusion and Decentralized Control in Autonomous Robotic Systems*, Pittsburgh, PA, Oct. 1997, SPIE, pp. 11–22.
- [8] J. Elson and D. Estrin, "Time synchronization for wireless sensor networks," in *IPDPS Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing*, Apr. 2001.
- [9] H.D. Young, *University Physics 8th Ed.*, Addison Wesley, 1992.
- [10] W.M. Siebert, *Circuits, Signals, and Systems*, MIT Press, 1995.
- [11] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Mobile networking for smart dust," in *5th ACM/IEEE MOBI-COM*. ACM/IEEE, Aug. 1999.
- [12] S. Hollar and K. Pister, "Toward networked micro robots," in *International Conference on Intelligent Robots and Systems*. IEEE/RSJ, Oct. 2001.
- [13] L. Girod, "Development and characterization of an acoustic rangefinder," Tech. Rep. CS-TR-00-728, University of Southern California, Mar. 2000.
- [14] D. Sarwate and M. Pursley, "Crosscorrelation properties of pseudorandom and related sequences," *Proceedings of the IEEE*, vol. 68, no. 5, pp. 593–619, May 1980.
- [15] P.R. Wolf and C.D. Ghilani, *Adjustment Computations: Statistics and Least Squares in Surveying and GIS*, Wiley, 1997.
- [16] O. Cramer, "The variation of the specific heat ratio and the speed of sound in air with temperature, pressure, humidity, and CO_2 concentration," *Journal of the Acoustical Society of America*, vol. 93, no. 5, pp. 2510–16, May 1993.
- [17] T. Kanade, M. Okutomi, and T. Nakahara, "A multiple-baseline stereo method," in *DARPA Image Understanding Workshop*. DARPA, 1992.