Institute for Software Integrated Systems
Vanderbilt University
Nashville Tennessee 37235

TECHNICAL REPORT

Acoustic Ranging in Resource Constrained Sensor Networks

János Sallai, György Balogh, Miklós Maróti, Ákos Lédeczi
Institute for Software Integrated Systems
Vanderbilt University, 2015 Terrace Place, Nashville, TN 37203, USA
{ sallai, bogyom, mmaroti, akos }@isis.vanderbilt.edu
http://www.isis.vanderbilt.edu

Fine-grained geographic localization of nodes is essential for an extensive range of distributed sensor applications. To compute geographic coordinates, localization algorithms commonly use pair-wise distance estimates between nodes. In this technical report we present a noise tolerant acoustic ranging mechanism for wireless sensors that employs digital signal processing techniques on standard MICA hardware. We describe how noise canceling, digital filtering and peak detection can be applied to meet the severe resource constraints of the platform, yet yielding average range estimation errors below 10cm independently from the actual node-to-node distances.

Keywords: Sensor Networks, Acoustic Ranging, Digital Signal Processing

1 Introduction

Wireless sensor networks consisting of small, low-power nodes equipped with different sensors and actuators have been gaining attention among researchers in the past few years. The fields of their possible applications range from military surveillance to precision agriculture. It is not uncommon that tolerance to severe environmental conditions, such as significant background noise or extreme temperatures, is a requirement. The inconvenience or infeasibility of human interaction in these scenarios raises a need for ad-hoc deployment and unattended operation.

As geographic location of nodes is required by a number of sensor applications and middleware services, such as positioning systems, collaborative sensing and signaling applications, and location-aware routing services, it is imperative that the sensor network be able to conduct self-localization. Wireless sensor networks are intrinsically different from traditional distributed systems due to the strict resource constraints on the sensor nodes. Resources are primarily constrained by energy consumption, hardware size and cost. System lifetime should be in the order of weeks or months, requiring low-power hardware as well as power-aware software solutions. The cumulative hardware cost of the system needs to stay low, even though the number of nodes employed in a particular real-world application can be large. Furthermore, application-specific hardware tends to be expensive due to the relatively high costs of design and manufacturing necessitating the usage of COTS hardware in large-scale sensor networks.

Localization in sensor networks is most commonly accomplished using range estimations between sensor nodes.[1] An extensive amount of research has been done into various ranging techniques in the past few years. If high accuracy was not considered the primary design criterion, received RF signal strength information (RSSI) and RF proximity based methods provide sufficient results [1] [2] [3]. The most effective techniques, which yield results sufficient enough to carry out fine-grained localization, however, are based on time of flight (TOF) measurements of signals.

Purely RF time of flight based techniques, such as GPS, have limited applicability in sensor networks, since they demand high precision measurements and synchronization. Acoustic signals have many

---

[1] Research has been done to investigating range-free localization approaches as well. See [17] for details.

advantages over RF based approaches. Since the sound propagates much slower in air than RF signals, TOF can be precisely estimated from the time difference of arrival (TDOA) of simultaneously emitted acoustic and radio signals. As opposed to RF based TOF measurement techniques, clocks on the nodes need not be explicitly synchronized, post-facto synchronization [4] suffices. Ultrasonic ranging techniques, such as described in [5] and [6] can attain higher precision than the ones using audible sound, however, they provide shorter effective range and require more expensive hardware.

The ranging mechanism presented in this paper uses acoustics and leverages the advantages described above. Unlike other implementations on the same hardware, which make use of the analog tone detector on the MICA sensor board, in our approach we sample the acoustic signals then digitally process it to estimate the time of flight. Processing includes reduction of Gaussian noise using multiple sampling, digital filtering, and detecting the offset of maximum energy in the resulting signal. Though this implementation is significantly more expensive than the ones using the tone detector with regards to memory requirements and computational costs, it is much less sensitive to background noise and has a longer effective range.

Though acoustic ranging augmented with digital signal processing has already been the subject of research within the scope of sensor networks, existing implementations target more heavyweight hardware (i.e. sensor nodes with PC-class capabilities). Our prototype is unique in a way that it targets severely resource constrained devices, equipped with 4 to 8 MHz microcontrollers and 4 kb RAM. After specifying the hardware requirements of the application in section 2, section 3 introduces our acoustic ranging approach. We present the digital signal processing techniques suitable for severely constrained hardware to carry out amplification and filtering, and explain how range estimates are computed from the recorded samples. Temperature dependence issues and calibration is discussed in section 4. Section 5 evaluates our experimental results; and Section 6 discusses the issues and limitations of our approach. Finally, we give a brief comparison between our approach and two existing acoustic ranging implementations in section 7.

2   Hardware

Our acoustic ranging application targets the MICA/MICA2 motes developed at UC Berkeley as a research platform for low-power wireless sensor networks [7].

The MICA mote is equipped with a 4 MHz RISC microcontroller, 4 kb RAM and a 916 MHz wireless transceiver capable of data transfer at 19.2 kbps with the radio range of 200 feet, and is powered by two AA batteries. The microcontroller has no support for floating point arithmetic or integer multiplications. The MICA2 mote has a more advanced microcontroller running at 7.3 MHz and its transceiver supports transfer rates up to 38.4 kbps with an increased radio range of 500 feet.

The basic sensor boards, compatible with both MICA and MICA2 motes, are equipped with a number of sensors and actuators. Among them, the microphone and the fixed-frequency sounder are utilized by the application introduced in this paper. The maximum attainable sampling rate is around 18 kHz; the nominal frequency of the sounder is 4.4 kHz.

3   Approach

The concept of acoustic ranging is based on measuring the time of flight of the sound signal between the signal source (also referred as the acoustic actuator, or simply actuator) and the acoustic sensor. The range estimate can be trivially calculated from the time measurement, assuming the speed of sound is known and is constant.

Employing a sophisticated synchronization mechanism is essential to accurately measure the time of flight. The most common approach is having the actuator notify the sensor via a radio message at the same time when the signal is emitted. Since the propagation speed of the radio signal is approximately 10^6 times higher than the speed of sound, the difference of the arrival times of the sound and radio signals is a good estimate of the time of flight in question.

However, there is a problem with the practical application of this approach, namely that it is the *start* of the signal that needs to be detected, which is cumbersome for the following reasons:

a. Generating a sound signal with a sharp rising envelope is infeasible with the available hardware.

b. Accurate detection of the start of a noisy signal is difficult.

To satisfactorily address this issue our ranging solution first computes the sample-wise sum of multiple sampled signals. This way the Gaussian noise in the original samples will cancel out, and the summed signal will have a better signal-to-noise ratio. Then, we apply a digital band-pass filter, and finally we detect the first peek in the filtered samples that will be used to estimate the start of the original signal.

3.1 Increasing the signal-to-noise ratio

To adequately address the problem of locating the beginning of the chirp, first we need to increase the signal-to-noise ratio of the samples.

In our approach, the acoustic signal consists of a series of chirps, all of the same length, with variable-length intervals of silence in between. Delays between the consecutive chirps are known to the sensor. Since the sensor knows the emission time of the series of signals (the sensor is notified via a radio message as discussed before) and the exact pattern as well, it can calculate the emission time of each chirp. The chirps are sampled one by one, then added together and processed as a single sampled signal.
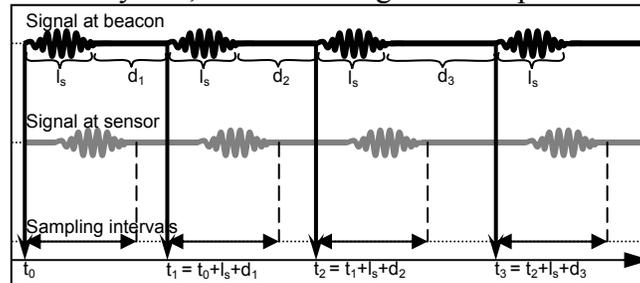

Figure 1. Sampling of multiple signals. The length of the signals (l_s) and the delays between the consecutive chirps ($d_1, d_2, \ldots, d_N$) are known to the sensor, this way the start times of the sampling intervals can easily computed.

Since disturbances such as ambient and electronic noise are of Gaussian nature, they are independent for each chirp, whereas the useful signal content will be identical. Adding together the series of samples improves the SNR by $10\lg(N)$ dB, where N is the number of chirps used. Our prototype uses 16 consecutive chirps in an acoustic ranging signal, thus the SNR is improved by 12 dB.

Delays between consecutive chirps are varied to avoid a situation when multiple samples have the same noise pattern at the same offset, which is a common phenomenon caused by acoustic multi-path effects. Hence the independent nature of the disturbances is preserved.

To keep the memory requirements at a minimum, our implementation uses an accumulator buffer for the sampled signals, where the additions are done on the fly.

3.2 Filtering

The acoustic signals are of a fixed frequency with slight variations between distinct actuator nodes, probably due to manufacturing differences. Lower and upper bounds for the frequencies were measured to be 4000 and 4500 Hz respectively; the sensors were, thus, tuned to search for the acoustic signals in that frequency range.

3.2.1 Designing the filter

To improve the SNR further, a digital bandpass filter is employed in our acoustic ranging mechanism. Since the ambient noise in our test recordings was found to be colored (with amplitude decreasing by 20 dB per decade below 2 kHz and approximately flat above) a matched bandpass filter was used.
The design criterion was primarily to increase the signal-to-noise ratio while keeping the integer filter coefficients in the [-4,4] interval and the tap number small to keep hardware requirements at a minimum. This way, calculation of a filtered sample can be accomplished using 4 accumulator variables, without multiplications, that would be compiled into additions on a processor that has no support for that.
The first accumulator variable is assigned to coefficients 1 and -1, the second to 2 and -2 and so on. In our prototype, for each tap, if the coefficient is positive we add the sampled value to the accumulator variable that corresponds to the filter coefficient. If the filter coefficient is negative, we do subtraction instead of addition. The total number of the above additions and subtractions is less than the tap number of the filter, since we do not have to do anything at the taps with 0 coefficients. Finally, we take the weighted sum of the accumulator variables[2] and then scale the result back with a binary shift.

3.2.2 Genetic search for the integer coefficients

There was a lot of research done to explore the applicability of evolutionary algorithms in digital filter design in the late nineties. The essential idea behind these approaches was to use evolutionary algorithms to optimize filter coefficients [8][9][10]. Though they were predominantly addressing hardware design issues, as [8] and [10], their problem domain has a lot in common with digital filter design for resource-constrained sensor network nodes. Consequently, the integer coefficients of the bandpass filter employed in our acoustic sensor application were calculated by a genetic algorithm.
In order to construct the fitness function for the genetic optimization algorithm, we recorded several windows containing both chirps and silence then applied the filtering to the signals in the way described before. The fitness function chosen was the signal-noise ratio, which can easily be estimated from the training signals, assuming that the positions of the chirps and the silence within the recordings are known.
The output of the genetic search was a 35-tap FIR filter with integer coefficients in the [-4,4] interval, which has a suppression of at least 12 dB below 3800 Hz and above 4500 Hz, and has a roll-off rate of approximately 20 dB per decade below 3800 Hz.
With the resulting tap number and coefficients we can calculate one filtered sample with 34 additions and subtractions and two shift operations.

3.3 Range estimation

The power of the filtered samples has a local maximum in the interval where a chirp is recorded. By detecting the peek of the signal power it is possible to give an estimate of the start of the signal.
Since calculation of power requires taking the squares of the samples, which is an expensive operation on a platform that does not support multiplication, we approximate the local maxima of the power

---

[2] This can be done by 5 additions and a binary shift: `weighted_sum = (a1 + a3) + ((a2 + a3 + a4 + a4) << 1`

function as follows. First we define a moving average function over the absolute value of the samples. Then we find the global average of the absolute value of the amplitude, so that later it will be possible to differentiate between signal and silence based on whether the value of the moving average function or the global average is higher at the given offset. Filtering, taking the absolute value, and averaging are carried out in the same loop in-place to minimize time and memory requirements.

Due to disturbances, even though the sample is filtered, it is possible that multiple local maxima of the moving average function are above the global average. We should, however find the local maximum that corresponds to the chirp, and discard all other noise patterns of significant energy that fall into the same frequency range.

For this reason, we examined the moving averages of the test samples around the positions of the chirps, and found that the moving averages of valid chirp patterns have segments with length of 200 to 350 samples above the average amplitude. Thus, we implemented the peek detection so that it returns the first local maximum that satisfies the above constraint. All other peeks are discarded.

4 Calibration

The distance between the actuator and the sensor is proportional to the time of flight of the acoustic signal. The peak detected, however, does not exactly reflect the time of flight, since it is obviously not the same offset that corresponds to the start of the acoustic signal, but some arbitrary one following that. The difference between the peak and the beginning of the signal is the result of the unknown rise time of the signal and the delay of the filter.

Consequently, before scaling the offset of the peak with a suitable constant (which is the number of distance units the sound travels during the time represented by one sample) to yield the range estimate, we need to compensate for this delay of various causes by an additive constant.

Since the latency in question is unknown, we chose to solve the problem statistically. A number of measurements were made with varying distances between sensor and actuator nodes then a linear regression was applied to the measured offsets of maximum energy and the actual distances. The additive and the multiplicative regression constants thus corresponded to the offset caused by the latency and the speed of sound respectively.

4.1 Dependence on the speed of sound

It is important to note here that the speed of sound in the air varies depending on the temperature, pressure, relative humidity as well as the CO_2 concentration of the air.

Using the formula given in [12] we can see that there is a 5% difference between the speed of sound at 0˚ C and at 30˚ C, which when used to calculate the range estimate from the time of flight of an acoustic signal would cause in half a meter error at a 10 meter distance.

The acoustic ranging component implemented does not compensate the ranging errors resulting from the variation of speed of sound depending on the environmental conditions. It assumes the speed of sound to be 340 m/s, and estimates the range accordingly. An application developer intending to use this component in an application could, for example, do one of the following:

a. *Compensate the ranging error based on thermometer reading* by multiplying the range estimate with the ratio of the hard-coded reference speed of sound of 340 m/s by that computed from the measured temperature.

b. *Compensate the ranging error based on location constraints.* Assume that a certain number of anchor positions (positions of actuator or sensor nodes) are known. If the acoustic ranging gives consistently less or greater distance estimates than the distances computable using the anchor positions, the ranging

estimates should be correspondingly scaled so that the relative coordinate system of range estimates match the absolute coordinate system defined by the anchor nodes.

c. *Accept the range estimates* without error compensation if the range estimation with the reference speed of sound is satisfactory, or there are no means of error compensation.

5   Results

We tested the acoustic ranging prototype with 50 MICA2 motes equipped with standard sensor boards. The test application consisted of the acoustic ranging component, a time slot negotiation component (to prevent two motes within each other's acoustic range from chirping at the same time), and middleware services such as routing and remote control. The application used only 3332 bytes of RAM. The ranging experiment was controlled from a PC, using a Java application that recorded the incoming rage estimates and could optionally carry out localization using a basic linear spring model.

The experiment was carried out in a parking lot. The air temperature at ground level was approximately 35° C with relative humidity of 60%. The motes were evenly distributed on a 15 by 30-meter area with no obstructions between any sensor pairs to assure direct line of sight. The actual distances were measured between the motes with an ultrasonic ranging device to enable the evaluation of the accuracy of the ranging approach.
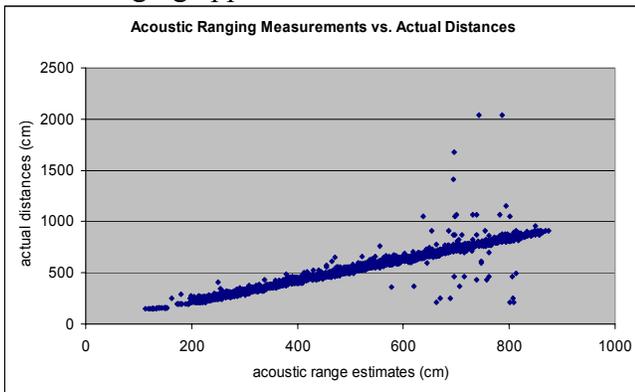

Figure 2. Acoustic range estimates vs. the actual distances. Outliers present due to a hardware problem of a single node.
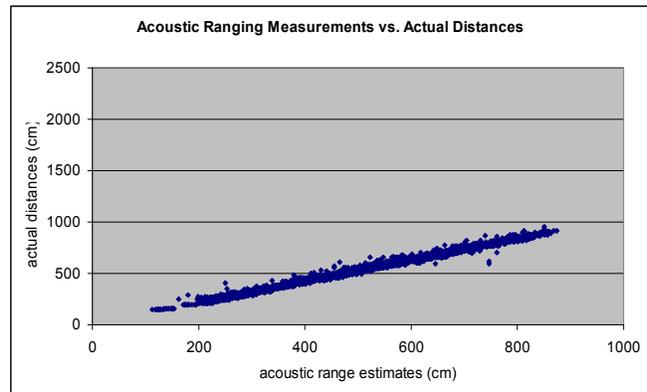

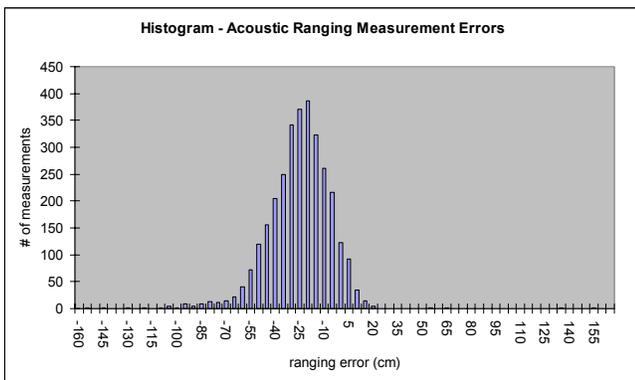**Figure 3.** Acoustic range estimates vs. the actual distances after removing outliers.

Figure 4. Histogram of acoustic measurement errors. Due to higher air temperature than the reference value the nodes underestimated the distances by 27.68 cm.
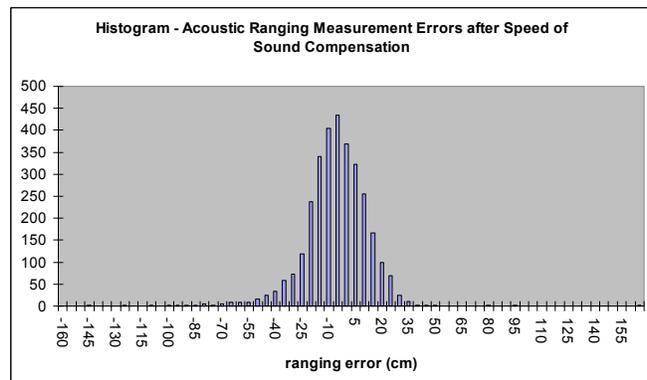

Figure 5. Histogram of acoustic measurement errors after compensating the difference between the calibrated and the actual speed of sound. Average error is -8.18 cm.

The acoustic ranging measurements were repeated ten times. Figure 2 shows the correspondence between the range estimates and the actual distances. As we can see, the relationship between them is approximately linear, with some random outliers. Analysis of the range estimates showed that most of the outliers were generated by a single, malfunctioning mote, so the corresponding measurements were removed not to disturb the further evaluation of the acoustic ranging and localization technique. Figure 3 shows the range estimates vs. the actual distances without the range estimates of the mote in question. Note that automatic elimination of these kinds of errors would be relatively simple.

As the histogram of the ranging errors (Figure 4) shows, the mean of the errors was around -28cm, that is, the motes underestimated the distances. This can be explained by the difference between the reference speed of sound used in the range estimator algorithm and the high actual speed of sound resulting from the relatively high air temperature. After adjusting the ranging estimates using the actual speed of sound, the average error is decreased to -8.18cm.

6   Issues and limitations

Fine-grained localization of low-powered, cheap nodes still eludes us after years of research in the domain of wireless sensor networks. There are inherent problems with acoustic ranging, such as their relatively limited range and the need to compensate measurement errors due to non-line-of-sight conditions.

6.1   Acoustic ranging errors

Generally, the error of acoustic distance estimation can be expressed as the sum of a Gaussian and a non-Gaussian component. The Gaussian component is the result of noisy measurements, the non-Gaussian part, on the other hand, is caused by multi-path effects.

While Gaussian measurement errors can be compensated successfully by averaging a series of consistent range estimates, the effects of echoes and obstructions cannot be adequately handled.

If the line of sight between the actuator and the sensor is obstructed, the sensor will consistently report a longer range estimate than the actual distance. In a purely acoustic localization system, the overall error caused by non-line-of-sight conditions can be mitigated through various heuristics (e.g. geometric consistency checks as described in [13]). However, building an entirely error-tolerant purely acoustic solution appears to be infeasible. As a possible way to improve the reliability of the self-localization, [14] suggests using multiple sensor modalities. A good example of such a technique is presented in [14], where the acoustic ranging mechanism is augmented with infrared LEDs and cameras to detect non-line-of-sight conditions.

6.2   Hardware limitations

The most serious hardware constraint of our acoustic ranging implementation is the limited availability of RAM. One sampled acoustic signal needs to fit into the buffer allocated for the acoustic ranging component. The size of the audio buffer was 1700 bytes in our application. The application consisted of the acoustic ranging component, the time slot negotiation algorithm and a multi-hop routing layer and the necessary TinyOS [15] system components.

An increase of the available memory would certainly increase the maximum range achievable by this technique. Yet it would still be possible to extend the effective range with the existing hardware by increasing the number of chirps in the acoustic signal and starting sampling them with increasing time delays. This way it is possible to explore a longer range with the same amount of available memory. However, it would take more time and consume more power.

Another possible approach would be to decrease the sampling rate when sampling the first series of chirps. The decreased sampling rate facilitates the exploration of a longer range with the same buffer size, obviously with decreased peek detection accuracy. In a next phase, however, another series of chirps could be sampled with a high sampling rate and with a suitable start delay to ensure that the chirps fall in the sampled time windows.

7   Comparison with existing acoustic ranging solutions

In the last few years there has been an abundance of publications on localization in sensor networks. However, they discuss mostly theoretical results; and only a fraction of them describe working prototypes. Below we contrast our solution with the acoustic ranging approach described in [13] and in [14], and the acoustic ranging mechanism underlying Calamari, the localization system presented in [16].

[13] and [14] presents an acoustic ranging system implemented on PC-class nodes equipped with a PC sound card. The acoustic signal emitted by the transmitter is formed by modulating a binary code using binary phase shift keying (BPSK) at a 12 KHz chip rate. The binary code is known to the detector, so it can compute the correlation between the reconstructed reference signal and the received signals at every possible offset to determine the position of the chirp. While this approach performs robustly, yielding distance estimates with sub-centimeter errors, it has a considerable computational complexity. In contrast, when designing our solution we were constrained by a fixed-frequency buzzer, a maximum sampling rate one third of that of a PC sound card and 4 kilobytes of precious RAM. The resource constraints forced us to apply simplified, less sophisticated signal processing mechanisms tailored to the given hardware; and as an agreeable tradeoff we were able to keep the average error of distance estimation below 10cm.

Calamari, the localization system introduced in [16], uses acoustic TOF-based distance estimations as the underlying ranging mechanism. The implementation targets the MICA platform; the motes are equipped with the standard MICA sensor board. Unlike our solution, Calamari uses the tone detector of the sensor board to identify the acoustic signal. Though using the analog hardware is cheaper than sampling and signal processing in all regards, its effective range is under 3 meters, and the uncalibrated distance estimates are very poor ([16] reports an average error of 74.6%). Applying sophisticated calibration methods in Calamari reduces the average error to 10.1%, however, the error, due to the use of the tone detector, is distance dependent. Our approach, though it consumes precious RAM and has some computational overhead, provides more accurate results with uniform errors within the effective range.

8   Conclusions

We have presented an acoustic ranging mechanism augmented by simple digital signal processing techniques that targets severely resource-constrained devices. We have increased the effective range of the acoustic distance measurements to nine meters with the average accuracy of 8cm on the MICA/MICA2 motes, which is a significant improvement over a ranging solution that relies purely on the analog tone detector of the sensor board. Even though digital signal processing usually implies computationally intensive tasks, which may seem rather expensive if used in low-power, resource-constrained sensors, our prototype implementation proved the viability of our approach.

9 Acknowledgements

10 References

[1]	N. Bulusu, J. Heidemann, D. Estrin, Gps-less low cost outdoor localization for very small devices. IEEE Personal Communications Magazine, Vol. 7, No. 5, p. 28-34, 2000

[2]	N. Bulusu, V. Bychkovskiy, D. Estrin, J. Heidemann, Scalable, Ad Hoc Deployable RF-based Localization. Grace Hopper Celebration of Women in Computing Conference 2002, Vancouver, British Columbia, Canada, 2000

[3]	P. Bergamo, G. Mazzini, Localization in Sensor Networks with Fading and Mobility. IEEE PIMRC, 2002.

[4]	J. Elson, D. Estrin, Time Synchronization for Wireless Sensor Networks. Proceedings of the 2001 International Parallel and Distributed Processing Symposium (IPDPS), Workshop on Parallel and Distributed Computing Issues in Wireless and Mobile Computing, San Francisco, California, USA. April 2001.

[5]	A. Ward, A. Jones, A. Hopper,A New Location Technique for the Active Office. IEEE Personal Communications, Vol. 4, No. 5, October 1997, pp 42-47.

[6]	N. Priyantha, A. Chakraborty, H. Balakrishnan, The Cricket Location-Support System. 6th ACM International Conference on Mobile Computing and Networking, 2000.

[7]	J. Hill, D. Culler, Mica: A Wireless Platform for Deeply Embedded Networks. IEEE Micro, Vol 22(6), (2002) 12-24.

[8]	Wade G., Roberts A., and Williams G., Multiplier-less FIR filter design using a genetic algorithm. IEE Proceedings in Vision, Image and Signal Processing, Vol. 141, No. 3, pp. 175180, 1994

[9]	D.J. Xu, M.L. Daley, Design of optimal digital filter using a parallel genetic algorithm. IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing, Vol. 42 , No. 10, p. 673-675, 1995

[10]	D. Quagliarella, J. Periaux, C. Poloni, G. Winter, Genetic Algorithms and Evolution Strategy in Engineering and Computer Science: Recent Advances and Industrial Applications. John Wiley & Sons, 1998

[11]	J. F. Miller, Digital Filter Design at Gate-level using Evolutionary Algorithms. Proceedings of the Genetic and Evolutionary Computation Conference, p. 1127-1134, 1999

[12]	O. Cramer, The variation of the specific heat ratio and the speed of sound in air with temperature, pressure, humidity, and CO2 concentration. Journal of the Acoustical Society of America, 93(5) p. 2510-2616, formula at p. 2514

[13]	L. Girod, D. Estrin, Robust Range Estimation for Localization in Ad-hoc Sensor Networks. 2000

[14]	L. Girod, D. Estrin, Robust Range Estimation Using Acoustic and Multimodal Sensing. International Conference on Intelligent Robots and Systems, 2001

[15]	TinyOS, web site at http://webs.cs.berkeley.edu/1/

[16]	K. Whitehouse, D. Culler, Calibration as Parameter Estimation in Sensor Networks. Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications, 2002

[17]	T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. F. Abdelzaher. Range-Free Localization Schemes in Large Scale Sensor Networks, MobiCom 2003.