

Countersniper System for Urban Warfare

ÁKOS LÉDECZI, ANDRÁS NÁDAS, PÉTER VÖLGYESI, GYÖRGY BALOGH,
BRANISLAV KUSY, JÁNOS SALLAI, GÁBOR PAP, SEBESTYÉN DÓRA, and
KÁROLY MOLNÁR

Institute for Software Integrated Systems, Vanderbilt University, USA

MIKLÓS MARÓTI

Bolyai Institute, University of Szeged, Hungary

and

GYULA SIMON

Department of Computer Science, Veszprém University, Hungary

An ad-hoc wireless sensor network-based system is presented that detects and accurately locates shooters even in urban environments. The localization accuracy of the system in open terrain is competitive with that of existing centralized countersniper systems. However, the presented sensor network-based solution surpasses the traditional approach because it can mitigate acoustic multipath effects prevalent in urban areas and it can also resolve multiple simultaneous shots. These unique characteristics of the system are made possible by employing novel sensor fusion techniques that utilize the spatial and temporal diversity of multiple detections. In this article, in addition to the overall system architecture, the middleware services and the unique sensor fusion algorithms are described. An analysis of the experimental data gathered during field trials at US military facilities is also presented.

Categories and Subject Descriptors: B.7.1 [**Integrated Circuits**]: Types and Design Styles—*Algorithms implemented in hardware*; C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing protocols*; G.1.0 [**Numerical Analysis**]: General—*Numerical algorithms*; J.7 [**Computers in Other Systems**]*—Military*

General Terms: Algorithms, Design, Measurement, Performance

Additional Key Words and Phrases: Sensor networks, middleware services, time synchronization, message routing, data fusion, acoustic source localization

This research was supported by the DARPA/IXO NEST program (F33615-01-C-1903).

Authors' address: Á. Lédeczi, A. Nádas, P. Völgyesi, G. Balogh, B. Kusy, J. Sallai, G. Pap, S. Dóra, and K. Molnár, Institute for Software Integrated Systems, Vanderbilt University, Box 1829, Station B, Nashville, TN 37235; M. Maróti, Bolyai Institute, University of Szeged, Szeged, Hungary; G. Simon, Department of Computer Science, Veszprém University, Hungary; Correspondence E-mail: akos.ledeczi@vanderbilt.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 1550-4859/05/1100-0153 \$5.00




Fig. 1. The sensor network based shooter localization system using shockwave and muzzle blast time of arrival data.

1. INTRODUCTION

Experiences from recent military conflicts clearly indicate that the armies of developed countries will increasingly face asymmetric warfare in the future. Countering this new threat will have to include the expansion of the technological advantage that modern armies already have. One example where technology plays an important role is countersniper systems, which can be extremely useful in urban guerilla warfare. However, current systems do not work well in urban areas because of the reverberant environment. Nor are they able to distinguish multiple simultaneous shots that are very common in firefights. A sensor network-based solution, however, can address both of these problems because of its widely distributed geographically diverse sensing capability.

Countersniper systems can use several different physical phenomena related to the shot or the weapon itself, such as acoustic, visual, or electromagnetic signals. A detectable visual event is the muzzle flash, as in case of the Viper system [Moroz et al. 1999], or the reflection from the sniper's scope [Vick et al. 2000]. The electromagnetic field or heat generated by the projectile can also be used for detection [Vick et al. 2000]. In spite of the wide range of possibilities, so far acoustic signals, such as the muzzle blast and the ballistic shockwave, provide the easiest and most accurate way to detect shots, and hence, the majority of existing countersniper systems use them as the primary information source [Duckworth et al. 1996; Pilar Systems; Stoughton 1997].

The most obvious acoustic event generated by the firing of any conventional weapon is the muzzle blast. The blast is a loud, characteristic noise originating from the end of the muzzle, and propagating spherically away at the speed of sound, making it ideal for localization purposes (see Figure 1). A less favorable property of the blast is that it can be suppressed by silencers or rendered ambiguous by acoustic propagation effects.

Typical rifles fire projectiles at supersonic velocities to increase both the range and accuracy, producing acoustic shocks along their trajectory. The shockwave is the result of the air being greatly compressed at the tip, and expanded at the end of the bullet, as it slices through the air. Under ideal circumstances the pressure signal detected by a microphone has a characteristic and distinctive waveform, called N-wave referring to its shape. Because of its very fast rise time ($<1 \mu s$), it cannot be produced by any other natural phenomenon. The ideal shockwave front is a cone (the Mach cone) moving along the trajectory of the projectile. The angle of the cone depends on the speed of the bullet. Note that this angle is continuously increasing as the bullet decelerates producing a distorted conical shape, as shown in Figure 1. Since N-waves can be accurately detected, shockwaves provide excellent means to determine projectile trajectories.

Existing commercial systems use a small number of microphone arrays (typically 1–3) to determine the projectile trajectory and/or the location of the shooter [Duckworth et al. 1996; Pilar Systems]. This centralized approach permits the use of practically any sophisticated signal processing algorithms working on the recorded acoustic signals of multiple microphones. On the other hand, the small number of sensor units limits the covered area and also makes the system sensitive to multipath effects and vulnerable to attacks. The proposed sensor networking approach allows the use of possibly several orders of magnitude higher number of inexpensive sensor units, but requires quite different processing approach because of the very limited communication bandwidth. Some of the processing must be allocated to the sensor units, while the sensor fusion needs to be carried out on a more powerful computer. The new technology presented in this article resulted in a robust, inexpensive, and highly accurate system, with some unique properties not available in any other systems, such as high insensitivity to multipath effects and the ability of handling multiple simultaneous shots. The concept is illustrated in Figure 1. The sensors accurately detect shockwave and/or muzzle blast events and measure their time of arrival (TOA). These timestamps of detected events are sent to a central base station, where the fusion algorithm calculates the shot trajectory and/or the shooter location, based on the TOA measurements and the known sensor locations. The communication in the network is provided by ad hoc routing protocols, incorporating the time synchronization service as well.

Muzzle blast and shockwave detections carry information about the shooter location, and the projectile trajectory, respectively. Either type of events or both combined can be used for localization purposes. The muzzle blast fusion algorithm works very well when the shooter is located within the sensor field and there are enough (at least 8–10) line-of-sight measurements. Once the shooter is shooting outside of the sensor field, the accuracy starts to decrease. One reason is that the angle of the sensor field from the shooter (“field of view”) is getting smaller and hence, individual measurement errors have larger effects on the result. The other reason is that, as the distance to the shooter increases, fewer and fewer sensors are able to detect the muzzle blast at all. Once the shooter is beyond 50 meters or so, muzzle blast alone is typically not enough to make accurate localization.




Fig. 2. The FPGA-based acoustic sensor board.

Shockwave fusion alone cannot determine the exact location of the shooter, but provides the trajectory of the bullet, even for long-range shots. The sensor network is presumably deployed in and around the protected area, and as long as the bullet goes through this region, the sensors can detect shockwave events, independently of the distance from the shooter. Naturally, shockwave trajectory estimation and muzzle blast ranging can be combined to provide accurate localization, if at least a few muzzle blast detections are available.

The outline of this article is as follows. The hardware and software architecture of the system is presented in Section 2. The shot detection component is described in Section 3. The routing integrated time synchronization is presented in Section 4. Fusion techniques (muzzle blast, shockwave, and combined) are presented in Section 5. The system is evaluated using field test measurements in Section 6.

2. ARCHITECTURE

2.1 Hardware Platform

The hardware platform is built upon the UC Berkeley MICA2 mote device running the TinyOS embedded operating system [Hill and Culler 2002], a widely used component-based architecture targeting wireless sensor network applications. Open interfaces at the software and hardware levels made it possible to integrate specialized smart sensor elements and supporting middleware services. Each MICA2 mote is furnished with an ATmega 128L 8-bit microcontroller with 128 kBytes instruction memory, 4 kBytes data memory and typical embedded peripherals built in. The on-board radio transceiver operates in the 433 MHz ISM band and has a maximum transfer rate of 38.4 kbits/sec with the maximum range of about 300 feet [XBOW].

Real-time detection, classification and correlation of acoustic events require processing power and buffer sizes not present in standard microcontroller-based embedded devices. To overcome these limitations, application-specific sensor boards have been designed and built at Vanderbilt University. The different architectures reflect the current dilemma faced by many signal processing engineers today.

The first version of the sensor board (see Figure 2) utilizes a Xilinx XC2S100 FPGA chip with three independent analog channels exploiting the inherent parallelism of the hardware. The algorithms—implemented in VHDL—are focusing




Fig. 3. The DSP-based acoustic sensor board.

on precise time domain analysis of acoustic signals captured at high sample rates (1 MSPS). Hardware and software interfaces (I²C bus, interrupts, led display and serial A/D) are implemented as custom IP cores in the same gate array. While this approach offers very appealing features, that is, high accuracy (note that on-board angle of arrival estimate is possible), high speed (though not fully utilized for audio purposes), and efficient resource utilization, the size of the FPGA component severely constrains the complexity of the applicable algorithms. Suboptimal power consumption of the processing unit and the lack of effective power management modes are other handicaps in the sensor network domain.

To overcome these limitations, another sensor board has been developed, where customized analog signal paths and an energy-efficient, powerful DSP processor make the unit uniquely suitable for power constrained applications. At the heart of the second platform (Figure 3) is a low-power fixed point ADSP-218x digital signal processor running at 50MHz. Its internal program (48KB) and data (56KB) memory buffers with advanced addressing modes and DMA controllers enable sophisticated signal processing and advanced power management methods.

Two independent analog input channels with low-cost electret microphones pick up the incoming acoustic signals utilizing 2-stage amplification with software programmable gain (0–54 dB). The A/D converters sample at up to 100 kSPS at 12-bit resolution. Analog comparators with software adjustable thresholds can be used to wake up the signal processor from low-power sleep mode, enabling continuous deployment for weeks on two AA batteries.

The FPGA and the DSP boards running the detection algorithms continuously draw 30 and 31 mA, respectively. In the power saving mode on the DSP board, this number drops to 1–5 mA, depending on the sleep mode. For comparison, the Mica2 mote draws 15 mA on average running the countersniper application.

2.2 Software Architecture

As the system evolved, different versions of the system architecture were described in detail in Simon et al. [2004] and Lédeczi et al. [2005]. Here, we present a summary of the latest software architecture (Figure 4). The Muzzle Blast and

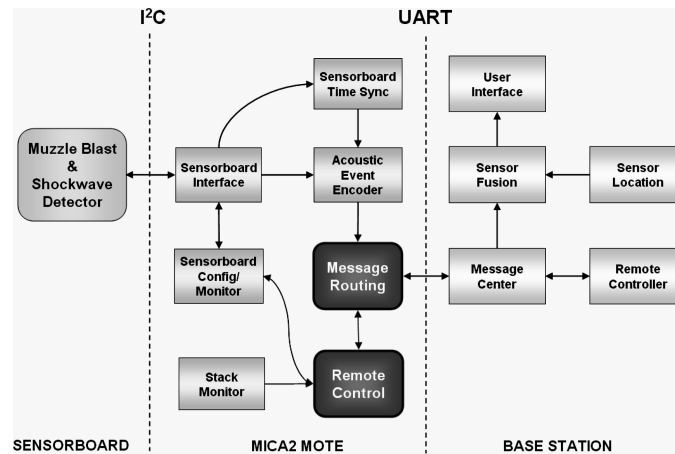


Fig. 4. The software architecture of the system.

Shockwave Detectors are implemented in VHDL on the FPGA of the first generation sensor board and in C on the DSP board (for details on detection refer to Section 3). The TOA data from either board is sent through the I²C interface to the mote. A separate software component translates the time from the clock of the sensor board to that of the mote. The Acoustic Event Encoder assembles a packet containing the TOA data and passes it to the Message Routing service.

In addition to transporting the packets to the base station through multiple hops, the Message Routing service also performs implicit time synchronization (see Section 4). Additional software components running on the mote include a Remote Control service enabling the configuration/polling of a single node, a group of or all of the nodes from the base station. A Stack Monitor makes sure that the limited memory of the mote is not exhausted.

The Base Station runs the Sensor Fusion algorithm utilizing the known sensor positions and displays the results on the User Interface. The accuracy and/or range of existing sensor self-localization methods (including our own [Sallai et al. 2004]) are not satisfactory for the shooter localization application. Hence, up-till-now all tests of the system were performed utilizing hand-placed motes on surveyed points. Future systems will use localization based on our accurate radio interferometric geolocation technique [Maróti et al. 2005].

3. DETECTION

A block diagram of the signal processing algorithm is shown in Figure 5. The incoming raw acoustic signal is compressed using zero-crossing (ZC) coding. The coded signals are used to detect possible occurrences of shockwave and muzzle blast patterns by the Shockwave Detection (SWD) and Muzzle Blast Detection (MBD) blocks, respectively. Although the operation of the two detection blocks is mainly independent, the SWD block can provide information (time of arrival of the detected shockwave) for the MBD block to facilitate the detection of a muzzle blast after a shockwave. Both blocks measure the TOA of the detected acoustic event using the on-board clock and then notify the mote. The Mica2




Fig. 5. Block diagram of the signal processing algorithm. The raw sampled signal is compressed and coded in the Zero-Crossing (ZC) Coder. The Shock Wave Detector (SWD) utilizes the ZC-coded signal, while the Muzzle Blast Detector (MBD) uses a filtered version. The detectors can communicate with the mote through an I²C interface.




Fig. 6. Zero crossing coding of the audio signal. Thin solid line shows the original signal, dashed lines are comparison values, and thick solid line represents the coded signal. In addition to the starting time (T), amplitude (Mm), length (L), and the rise time (τ) shown on the plot, the ZC code also contains the previous average amplitude (P) values.

mote reads the measurement data (TOA and optionally signal characteristics) and performs time synchronization between its own clock and that of the acoustic board. The measurement data is then propagated back to the base station using middleware services of the sensor network.

The zero-crossing coding is illustrated in Figure 6. Using a comparator with a deadband around zero, the intervals between zero-crossings are coded by storing the start time of the interval (T), length of the interval (L), the minimum or maximum signal value (Mm), the previous signal average amplitude (P) and the rise time (τ) of the signal. The SWD block utilizes the ZC-code stream to detect occurrences of shockwaves, while the MBD uses a filtered version of the same data-stream to find muzzle blast patterns. The ZC-filter block performs a nonlinear prefiltering operation, implemented very efficiently in the ZC-domain, to aid the muzzle blast detection.

Both shockwave and muzzle blast events are individually modeled by state machines, where states represent warped time and the transitions are governed by ZC properties [Simon et al. 2004]. The state machines were tuned using an extensive acoustic library of shots, converted to the ZC domain.

The signal detection algorithm proved to be quite robust. It recognized 100% of the training events and more than 90% of the other recorded shot events. (Note that, in reality, a shot may be detected by some sensors and may not be recognized by others, depending on the location of the sensor.) False positives could be produced only by physical contact with the microphone itself [Simon et al. 2004].

4. ROUTING INTEGRATED TIME SYNCHRONIZATION

Time synchronization is essential in the sensor network to make the distributed TOA measurements meaningful. Explicit time synchronization methods, for example, Elson et al. [2002], Ganeriwal et al. [2003], and Maróti et al. [2004], require periodic (re)synchronization of the nodes through the communication channel, thus increasing network traffic and power consumption. Early versions of the countersniper system [Simon et al. 2004] also used an explicit time synchronization protocol described in Maróti et al. [2004], producing very accurate synchronization.

As it is pointed out in Huang et al. [2000], post facto synchronization is enough in many cases, thus no continuous synchronization is required. Systems collecting data or reacting to rare events, but requiring exact time measurements belong to this class of applications. A post facto synchronization approach described in Huang et al. [2000] utilizes explicit pair-wise synchronization after message passing. An alternative method [Kusy et al. 2005] embeds the synchronization into the message routing protocol. This solution does not require any additional message exchange apart from the routing messages, but does require precise message time-stamping on both the transmitter and receiver sides. Later versions of the countersniper system utilized this approach [Lédeczi et al. 2005].

When a sensor in the network detects an event, it is time-stamped using the sensor's local, unsynchronized clock. The target node, that is, the base station, possibly several hops away from the sensor, needs to know the time of the event in its own local time. Without explicit synchronization in the network, the routing process can be used to perform implicit time synchronization.

Along with the sensor reading, a radio message includes an *age* field, which contains the elapsed time since the occurrence of the acoustic event. This additional information adds only a very small overhead to the message. Each intermediate mote measures the offset, which is the elapsed time from the reception of the packet with the sensor reading till its retransmission. The *age* field is updated upon transmission using a precise low-layer time stamping method described in Maróti et al. [2004]. When the sensor reading arrives at the destination, the *age* field contains the sum of the offsets measured by each of the motes along the path. The destination node can determine the time of the event by subtracting *age* from the time of arrival of the message. The concept is illustrated in Figure 7. An event is detected at node A at time instant T_{EVENT} , then a notification message is sent to destination node S through nodes B and C. The message delays at the nodes are offset_A , offset_B , and offset_C , respectively. The message arrives at S at time instant T_{rcvS} containing an age field of $\text{offset}_A + \text{offset}_B + \text{offset}_C$. The time of the event can be calculated as $T_{\text{EVENT}} = T_{\text{rcvS}} - \text{age}$.

The above method assumes that the offset values measured at different nodes use the same time unit. In reality, however, clock frequencies differ slightly. Since skew errors are not compensated for, error can accumulate during the routing process. The accuracy of the clock used in the Mica2 mote is better than 50 ppm, thus the worst-case post-facto synchronization error can be estimated as $5 * 10^{-5} T_R$, where T_R is the longest possible time of the message routing.

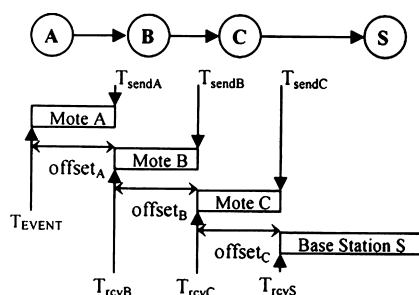


Fig. 7. Estimation of detection time T_{EVENT} can be iteratively determined along a routing path A, B, C, S as $T_{rcvS} - \text{offset}_C - \text{offset}_B - \text{offset}_A$. The horizontal axis is the timeline.

In the case of the countersniper system, messages older than two seconds have no value, thus the max synchronization error is $100 \mu s$. Experiments similar to the typical countersniper application setup showed that the synchronization error in a network of diameter 10 containing 45 nodes was $8 \mu s$ on average, while the largest pairwise error observed was $80 \mu s$ [Kusy et al. 2005].

This time synchronization protocol can be further refined by exploiting advantageous properties of certain wireless routing protocols, for example, messages arrive at the base station multiple times along possibly different paths [Maróti 2004]. Statistical analysis of multiple readings provides better estimate on the time the event occurred [Lédeczi et al. 2005].

The uniquely advantageous features of the proposed integrated time synchronization and routing algorithm, that is, no additional radio messages, support for power management, and the low imposed overhead on message size, make the algorithm suitable for many power-aware data collecting applications [Kusy et al. 2005].

5. SENSOR FUSION

The TOA measurements originating from either the muzzle blast or the shockwave can be used in the estimation process. Muzzle blasts are extremely useful in near-field position estimation [Simon et al. 2004], while shockwaves provide effective means to determine the direction of a distant shooter [Balogh et al. 2005]. In the next sections both methods will be reviewed and evaluated.

5.1 Muzzle Blast Fusion

There are a multitude of techniques for locating a static transmitting source by an array of listening devices. Near-field beam forming methods are successfully used to detect multiple sources in noisy reverberant areas [Chen et al. 2002, 2003]. However, the most sophisticated methods require the transmission of the sampled data records between nodes and/or the base station, for example, Chen et al. [2002]. A typical sensor network does not have the necessary communication bandwidth to support this alternative. There exist similar two-step techniques where in the first step the TOA, or equivalently the time difference of arrival (TDOA) data is calculated (or alternatively, measured),

and in the second step the location is calculated [Chen et al. 2002; Huang et al. 2000; Mahajan and Walworth 2001]. The communication burden of transmitting measured TOA or TDOA data is acceptable.

Since a pair of sensor readings defines a hyperboloid surface in space, in theory, four noncoplanar measurements are enough to determine a 3D location, provided the speed of sound is known. Unfortunately, errors in detection, sensor localization, and time synchronization all affect the accuracy of the solution. Using more measurements and solving the over-determined equations help overcome this problem [Chen et al. 2003; Huang et al. 2000; Mahajan and Walworth 2001].

Conventional methods (e.g., ones using LS or maximum likelihood criteria) work fine with noisy or even reverberant data, but in many cases, sensors not having direct line-of-sight detect echoes *only*, resulting in large errors in localization. In our experiments, in urban areas, typically 10–50% of the sensor readings provide erroneous data. Unfortunately, published localization methods do not address the problem of incorrect (TOA or TDOA) measurements. Simply applying the analytical solution or any other solution technique using the whole data set possibly containing a large number of incorrect measurements is not an option when high accuracy is required. Searching for the maximal set of consistent measurement data by repeatedly applying the solver on different subsets of input data is a possible solution, but not always computationally feasible.

The solutions proposed in Simon et al. [2004], and later refined to handle multiple shots in Lédeczi et al. [2005], utilized time of arrival data of the measured muzzle blasts. (Note that shockwaves in the early versions of the system were used to determine the trajectory of the projectile *after* the shooter location was determined using muzzle blasts [Simon et al. 2004].)

To find the position of the shooter(s), a *consistency function* on the four-dimensional space-time space is defined. A fast multiresolution search algorithm recursively finds the maxima of this function, which correspond to the location and time of possible shots. Then, these maxima are further analyzed to eliminate false positives caused by echoes. The consistency function is defined in such a way that it automatically classifies and eliminates erroneous measurements and multipath effects.

5.1.1 Consistency Function. Let N be the number of TOA muzzle blast measurements, and for each $i = 1, \dots, N$ let (x_i, y_i, z_i) be the coordinates of the sensor making the i th measurement, and t_i the time of arrival of the detected muzzle blast. Note that the measurements can contain good (line-of-sight) and bad (delayed-echo) readings. Also, since, in an urban combat scenario, several shots can be fired in a few seconds, the measurements may either correspond to a single or multiple shots as well, not known a priori to the algorithm.

For any hypothetical shooter position (x, y, z) and shot time t , the theoretical time of arrival of the muzzle blast at the sensor that recorded the i th measurement is

$$t_i(x, y, z, t) = t + \frac{\sqrt{(x - x_i)^2 + (y - y_i)^2 + (z - z_i)^2}}{v_{\text{sound}}},$$

where v_{sound} is the speed of sound. If the i th measurement is a direct line-of-sight detection of this hypothetical shot, then, in theory, the times $t_i(x, y, z, t)$ and t_i must be equal. In practice, however, due to errors in sensor localization, time synchronization and signal detection, only the following inequality is satisfied

$$|t_i(x, y, z, t) - t_i| \leq \tau, \quad (1)$$

where $\tau = \delta_1/v_{sound} + \tau_2 + \tau_3$ is an *uncertainty value*, δ_1 is the maximum sensor localization error, τ_2 is the maximum time synchronization error, and τ_3 is the maximum allowed signal detection uncertainty. For practical purposes, the localization error dominates τ . We assume that an upper bound for τ is known based on an a priori evaluation of the sensor localization, time synchronization and signal detection algorithms. The consistency function $C_\tau(x, y, z, t)$ is defined as the number of measurements for which (1) holds:

$$C_\tau(x, y, z, t) = \text{count}_{i=1, \dots, N} (|t_i(x, y, z, t) - t_i| \leq \tau).$$

The value of the consistency function for any (x, y, z, t) defines the number of measurements supporting the hypothesis that the shot was taken from (x, y, z) at time t , with uncertainty value τ . The consistency function is integer valued and always less than or equal to N . It is additive for the list of TOA measurements, and increasing in τ . Although $C_\tau(x, y, z, t)$ is not continuous, it satisfies the crucial property utilized in the discrete search algorithm:

$$C_{\tau/4}(x', y', z', t') \leq C_\tau(x, y, z, t), \quad (2)$$

whenever $|x - x'|/v_{sound}, |y - y'|/v_{sound}, |z - z'|/v_{sound}, |t - t'| \leq \tau/2$.

The consistency function usually takes its maximum value not in a single point, but in a 4-dimensional area, called the *max area*. The size of the max area depends on τ . If erroneous measurements are present, it is theoretically possible that there are more unconnected max areas, and it is also possible that the true location of the shot is not contained in the max area. Although simple counterexamples can be easily generated using enough bad measurements, in practical situations local maxima generated by multipath effects (mirror effect) are smaller than the maximum corresponding to the true shooter position [Simon et al. 2004].

The countersniper system recursively calculates the maximum of the consistency function as the location (and time) estimate of the shot. Since gradient-type search methods do not guarantee global convergence on a surface with multiple local maxima, a different method is utilized. Our fast multiresolution search algorithm finds the global maximum by searching the relevant sections of the space-time only, iteratively zooming to the correct spot.

5.1.2 Search Algorithm. The time complexity of finding the maximum of the consistency function in the guarded area $[X_{\min}, X_{\max}] \times [Y_{\min}, Y_{\max}] \times [Z_{\min}, Z_{\max}]$ and in the appropriate time window $[T_{\min}, T_{\max}]$ using a simple search is linear in terms of $X_{\max} - X_{\min}, Y_{\max} - Y_{\min}, Z_{\max} - Z_{\min}, T_{\max} - T_{\min}$ and N , because by (2) it is enough to evaluate $C_{\tau/4}(x', y', z', t')$ at grid points of the search space with uniform distance $v_{sound}\tau/2$ for the x, y and z coordinates,

and $\tau/2$ for t , and then finding the maximum among these points. However, the number of computation steps quickly becomes astronomical in practice, exceeding 10^{12} in our field experiments, rendering this simple algorithm not viable.

There is an extensive literature on several well-known algorithms for finding the local and global maxima of nonlinear functions, such as the Newton, Levenberg–Marquardt and Generalized Bisection methods. Since the consistency function is not continuous and we are interested in finding its global maxima, we applied the Generalized Bisection method based on interval arithmetic [Xu and Yang 2002]. Interval arithmetic introduces algebraic operations on closed intervals that represent possible values of variables. Every algebraic expression, including our definition of the consistency function, can be evaluated for intervals. For intervals $[x_{\min}, x_{\max}]$, $[y_{\min}, y_{\max}]$, $[z_{\min}, z_{\max}]$ and $[t_{\min}, t_{\max}]$, the consistency function yields the interval

$$[C_{\tau}^{\min}, C_{\tau}^{\max}] = C_{\tau}([x_{\min}, x_{\max}], [y_{\min}, y_{\max}], [z_{\min}, z_{\max}], [t_{\min}, t_{\max}])$$

that have the property

$$C_{\tau}^{\min} \leq C_{\tau}(x, y, z, t) \leq C_{\tau}^{\max}$$

for every $x_{\min} \leq x \leq x_{\max}$, $y_{\min} \leq y \leq y_{\max}$, $z_{\min} \leq z \leq z_{\max}$ and $t_{\min} \leq t \leq t_{\max}$.

The value C_{τ}^{\min} is the number of measurements that satisfy (1) for all points of the 4-dimensional rectangular region determined by $[x_{\min}, x_{\max}] \times \dots \times [t_{\min}, t_{\max}]$, while C_{τ}^{\max} is the number of measurements that satisfy (1) for some point of the same region.

During the search, we maintain a list of 4-dimensional rectangular regions (“boxes”), initially containing only $[X_{\min}, X_{\max}] \times [Y_{\min}, Y_{\max}] \times [Z_{\min}, Z_{\max}] \times [T_{\min}, T_{\max}]$, together with their evaluation under the consistency function. At each step, we remove the region that has maximum C_{τ}^{\max} value from the list, bisect it into two equal parts along its longest dimension, and insert the two resulting regions back to the list. We stop this procedure when the size of the maximum region is less than $v_{\text{sound}} \tau/2$ for the space and $\tau/2$ for the time coordinate. The resulting 4-dimensional region is guaranteed to contain at least one global maximum point of the consistency function $C_{\tau}(x, y, z, t)$. Note that there may be several boxes with the same C_{τ}^{\max} value, usually covering a small area around the true location. When displayed, this area provides an easily understandable visual representation of the uncertainty region of the location estimate.

When the global maximum with consistency value K is found, the K measurements contributing to it are removed and a new consistency function with the rest of the measurements is defined, and the search algorithm is restarted, until sufficiently high peaks (in practice with consistency value greater than 7–8) are detected. Note that the removal of “used” measurements speeds up the next search rounds as well.

Subsequent maxima correspond to either separate shots or echoes. An echo detection algorithm removes candidates possibly caused by multipath effects. The elimination process is based on the simple phenomenon that for a true

position (x, y, z, t) and its mirror image (x_M, y_M, z_M, t_M) it is true that $t \approx t_M$. If multiple peaks correspond to the same shot time within a limit (a few times $\tau/2$), only the one with the highest consistency value is kept, the others are labeled as echoes. Mirror images are formed by reverberated signals. Since the surfaces of buildings, other structures, etc. do not behave as one perfect mirror, the consistency value of mirror images in practice is seldom high, and almost always smaller than that of the true location, where good quality line-of-sight measurements are used. Of course, the method can be fooled if sensors are placed so that more sensors can sense the same reverberated signal than the line-of-sight signal. It is also possible that some of the true shots are eliminated as echoes, if their shot times are too close to each other. However, these are all unlikely events.

The speed of the search algorithm is remarkable. Compared to the simple linear search, the proposed solution provides more than eight orders of magnitude smaller computational complexity, resulting in approximately one second latency in field tests [Simon et al. 2004].

5.2 Shockwave Fusion

As described in Section 1, the shape of the Mach-cone is constantly changing as the bullet decelerates. However, we assume constant bullet speed within the sensor network to simplify the necessary calculations. The approximation is valid if the size of the sensor field is within a hundred meters, as shown by the remarkably good results in Section 6.4. In case of larger deployments, the fusion procedure can restrict the set of sensor readings considered so that they satisfy this constraint.

Sensors can identify the shockwave based on its unique waveform and the significant energy level. Based on the shockwave TOA measurements the bullet trajectory can be computed. The problem, similarly to the muzzle blast-based localization, is exacerbated by noisy and erroneous measurements, and possibly multiple simultaneous shots. Thus, a numerical approach, based on minimizing an error function, is applied.

The consistency function-based technique utilized in the muzzle blast fusion could easily be generalized by combining muzzle blast and shockwave measurements to contribute to the same consistency function. However, the number of dimensions would increase significantly. In addition to the shooter position (x, y, z) and the shot time t , the two angles for azimuth and elevation, and the speed of the projectile would result in a 7-dimensional space, where the search is currently computationally infeasible. Instead, a genetic algorithm was applied for shockwave fusion.

The real-time fusion algorithm receives shockwave measurements from the sensor network. First, the slicing algorithm splits the received measurements into smaller groups representing single shots or possibly almost simultaneous multiple shots. Then a genetic algorithm searches for the trajectories in each group that match the measurements the best. Once the best trajectories have been found and muzzle blast measurements are available, the range is also estimated.

5.2.1 Slicing. The slicing algorithm is responsible for splitting the incoming measurements into relatively small, processable groups. The slicing tries to provide groups with small size, ensuring that measurements from obviously different shockwaves do not belong to the same group.

A rough group formation is done by bounding the max time difference between any two measurements in the group. Then, the shockwave measurements in a group are sorted by time and split into smaller groups if and where there are larger time gaps than D/v_{sound} , where D is the maximum of pairwise distances between sensors in the field. The rationale behind this step is that two shockwave measurements with a larger time difference cannot be the result of a single shockwave because the wave front travels at the speed of sound and the projectile travels even faster [Balogh et al. 2005].

5.2.2 Trajectory Search. Given n shockwave measurements:

$$s_1(x_1, y_1, z_1, t_1), s_2(x_2, y_2, z_2, t_2), \dots, s_n(x_n, y_n, z_n, t_n),$$

where x_i, y_i, z_i are the coordinates of sensor i , and t_i is the TOA of the shockwave at sensor i , we are looking for multiple trajectories $tr_1(X_1, Y_1, Z_1, \alpha_1, \beta_1, v_1), \dots, tr_m(X_m, Y_m, Z_m, \alpha_m, \beta_m, v_m)$ that can generate the given measurements, where X_j, Y_j, Z_j is a point of the trajectory, α_j is the azimuth, β_j is the elevation and v_j is the speed of the projectile. The X_j, Y_j, Z_j point of the trajectory is the point where the bullet was at time t_0 , where $t_0 = 1/n \sum_{i=1}^n t_i$. This is a somewhat arbitrary choice for a single point on the trajectory that is close to the sensor field.

For a given trajectory $tr_j(X_j, Y_j, Z_j, \alpha_j, \beta_j, v_j)$ and a shockwave measurement $s_i(x_i, y_i, z_i, t_i)$, the theoretical time of arrival of the shockwave t'_i at the sensor can be calculated using simple Euclidean geometry.

From the measured and the theoretical time of arrivals, an error function of the trajectory is calculated. A genetic algorithm (GA) has been used to find the trajectory with the smallest error. The brief description of the GA is the following:

- (1) Generate an initial population of Q random trajectories.
- (2) Select w individuals randomly from the population.
- (3) Evaluate each individual in the selected subset using the *error function* (described later).
- (4) Sort the subset according to error.
- (5) Remove the worst 20% of the individuals in the subset, then generate new individuals by selecting random parents from the best 20% and applying genetic operators on the parents.
- (6) Go to Step (2).

Typical parameters used were $Q = 5000$ and $w = 500$. A general problem with GA is that it can get stuck in a local minimum when the whole population becomes homogeneous containing the same exact individuals. Different heuristics have been proposed to avoid this problem. We use the tournament selection technique [Blickle and Thiele 1995] to give a chance

for individuals with higher error value to breed and to slow down the homogenization.

In the presence of multiple shots or erroneous measurements (i.e., echoes), the whole set of measurements cannot be used: nonecho measurements belonging to the same shot need to be identified. Therefore, we have extended the representation of a trajectory with the subset of the measurement indexes. This way not only does the GA search for the trajectory, but also for a set of consistent measurements for which the trajectory has the smallest error. Once the best trajectory is found, the corresponding events are removed from the group and the search is started again using the rest of the data. It is repeated until the number of events becomes too small to be able to define a new solution. Since the trajectory has six parameters, at least six measurements are necessary.

The representation of a possible solution is the following: $sol(x, y, z, \alpha, \beta, v, S)$, where $x, y, z, \alpha, \beta, v$ is the trajectory parameters and S is a subset of shockwave measurement indices from 1 to n .

The following error function has been used to evaluate a given solution $sol(x, y, z, \alpha, \beta, v, S)$:

$$error = \begin{cases} \frac{1}{|S|} \sqrt{\sum_{i \in S} (t_i - t'_i)^2} & |S| \geq K \\ \infty & |S| < K, \end{cases}$$

where $i \in S'$ iff $i \in S$ or $|t_i - t'| < T$.

The error function is infinite if the number of selected measurements is smaller than K . Otherwise, an extended set of measurements S' is selected to calculate the error, containing elements of S and also the ones that has smaller time error than T . The value of K was set to 7 after experimentation with field data, while T was set to the value of the largest possible measurement error resulting from the sensor localization error [Balogh et al. 2005].

5.3 Range Estimation

Using the estimated trajectory based upon shockwave measurements and at least a few muzzle blast detections, the range of the source can also be estimated. Again, special care must be taken of potential multipath measurements and multiple shots.

Once the trajectory is estimated, the projectile location (X, Y, Z) at time instant t_0 is available, along with the elevation, azimuth and the speed of the projectile. Thus, the complete timeline of the bullet can be computed, provided the speed of the projectile is constant, by associating a time instant t with each bullet location (x, y, z) on the trajectory. As the unknown location of the shooter is also on the trajectory, a straightforward solution is to correlate the timeline of the bullet and the muzzleblast TOA data, using a simplified consistency function approach. Note that the search in his case is reduced to the one-dimensional space along the estimated trajectory, as follows:

A sliding window is moved backward on the trajectory to find the shooter position. The simplified consistency function at the trajectory position (x, y, z)

and its corresponding time t is defined as

$$C_\delta(x, y, z, t) = \text{count}_{i=1, \dots, N} (d_{\min}^i < (t_i - t)v_{\text{sound}} < d_{\max}^i),$$

where d_{\min}^i and d_{\max}^i are the minimum and maximum distances, respectively, between sensor position (x_i, y_i, z_i) (with TOA measurement t_i) and the sliding window of width δ , centered at (x, y, z) along the trajectory. The width of the window is determined by the estimated detection errors, a typical number being 1 meter.

The window position (x, y, z) with the highest consistency value $C_\delta(x, y, z, t)$ gives the estimated origin of the shot and, hence, the range. Similarly to the muzzle blast case, this solution automatically eliminates the erroneous measurements, and also the measurements corresponding to other shots.

6. EVALUATION

During the development period, several field tests were conducted in two U.S. Army facilities to evaluate the performance and accuracy of the shooter localization system. The data collected in the field tests were used to determine the accuracy of the system and its sensitivity to various sources of errors. In the following sections, fusion technologies based on the muzzle blast and the shockwave are analyzed separately.

6.1 Error Sources and Analysis Techniques

The sensor fusion algorithms use TOA measurements recorded by different sensors at different known locations. Hence, potential sources of measurement error are imperfect time synchronization and inaccurate sensor locations. The accuracy of the localization is also affected by the number and density of sensors used. Naturally, the acoustic environment also has a great impact on the performance of the system. In our experiments, urban test facilities were used, providing challenging environments rich in multipath effects.

In the field trials, the system was operated to provide its best performance, thus the maximum number of available nodes were utilized with precise time synchronization service, and the sensor localization was performed with approximately 20–30 cm precision. The data gathered at the field trials enabled us to experiment with the effects of potential error sources on the overall system accuracy offline. To perform the sensitivity analysis, the measured TOA values and the sensor locations were randomly modified to simulate the effects of additional time synchronization and sensor localization errors. Also, readings of randomly selected sensor units were neglected to study the effect of decreasing sensor density.

6.2 Muzzleblast Fusion Results

One of the field tests was performed in the MOUT training facility in Ft. Benning, GA, in July 2003. In this setup, 56 nodes were deployed in the central area of the McKenna village as shown in Figure 8, a screen dump of the system



Fig. 8. 2D system display, showing the aerial photo of the test site.

graphical user interface that includes an overhead picture of the MOUT site. The estimated position of the shooter is shown by the large circle, while the direction of the shot is indicated by an arrow. (Note that the shot direction was estimated from shockwave data, using an alternative approach described in Simon et al. [2004].) Other circles indicate the sensor positions where medium-sized ones denote sensor locations whose data were utilized in the current localization estimation.

The sensors were hand-placed, with an estimated position error of 20–30 cm. For error analysis purposes, 20 different known shooter positions were used in the experiment. During the test, 171 shots were fired, 101 of which were blanks and 70 were short range training ammunitions (SRTAs). Since the performance of the system was similar for both types of ammunition, only the unified results are presented.

The localization error of the system is shown in Figure 9, where the 3D error is the total localization error, while in the 2D error the elevation information is omitted. The system accuracy is remarkably good in 2D: the average error was 0.6 m, 83% of shots had less than one meter, and 98% had less than 2 meters of error.

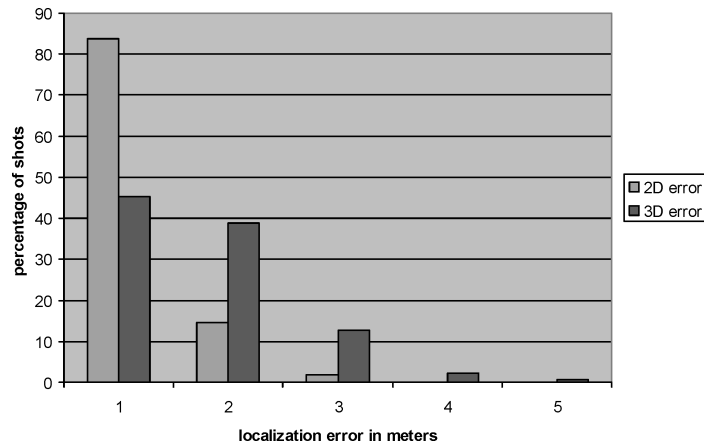


Fig. 9. Histogram of localization accuracy in 3D and 2D.

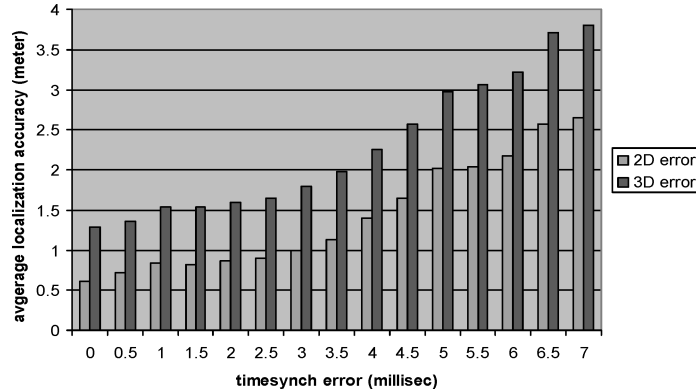


Fig. 10. Localization accuracy vs. time synchron error.

In 3D, the average error was 1.3 m, 46% of the shots had less than 1 m, and 84% of shots had less than 2 m of localization error.

6.3 Muzzleblast Fusion Sensitivity

The effects of time synchronization error are summarized in Figure 10. For each simulated time synchronization error value of T , the detection time for each sensor was modified by a uniformly distributed random value Δt , where $-T/2 < \Delta t < T/2$. Then the sensor fusion algorithm was run to estimate the shooter position. Each shot was used ten times; therefore, each data point in the diagram represents 1710 experiments.

The results in Figure 10 clearly show that the time synchronization accuracy is much better than what is needed by this application. The added 3D localization error of 10 cm in the presence of 0.5 ms time synchronization error is insignificant. On the other hand, for multiple shot detection and echo discrimination, well-synchronized measurements are advantageous.

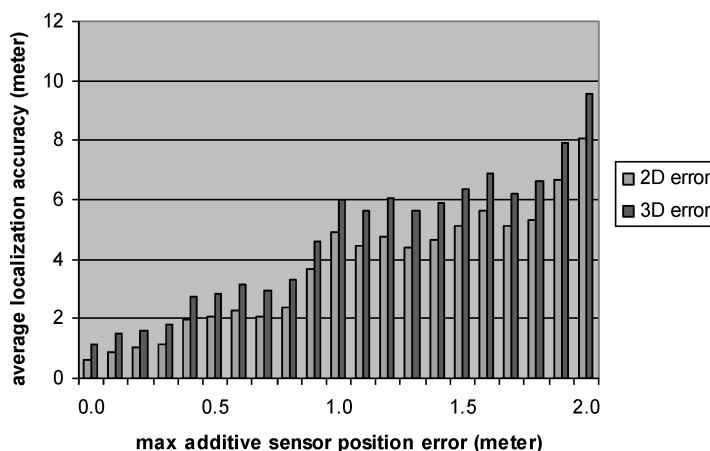


Fig. 11. Localization accuracy vs. sensor position error.

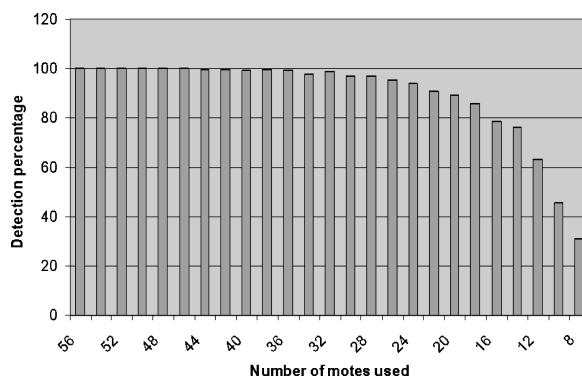


Fig. 12. Detection rate vs. number of sensors used.

The effect of the sensor localization error was simulated in a similar manner. The sensor positions were altered by adding uniformly distributed random values Δ , $-\Delta_{\max} < \Delta < \Delta_{\max}$, to the true positions originally measured and used in the field experiments. The fusion algorithm was run with the original measured TOA values, but used the modified sensor position data. Note that the “true” sensor positions had approximately 0.3-m position error, which is not included in Δ_{\max} . The results are shown in Figure 11, where 1000 randomized runs were performed and averaged for each Δ_{\max} value. The localization error is a nicely bounded, approximately linear function of the position error for small Δ_{\max} values. It’s clear from the results that in practice the sensor localization error dominates the negligible time synchronization error.

We have also analyzed the effects of sensor density. Again, we used the real data gathered on the field and then removed sensors randomly. The results are shown in Figure 12 and Figure 13. For each N , where N is the number of sensors and $N \leq 56$, we generated a random subset of the 56 available nodes, ran the sensor fusion for all the shots and repeated the procedure ten times.

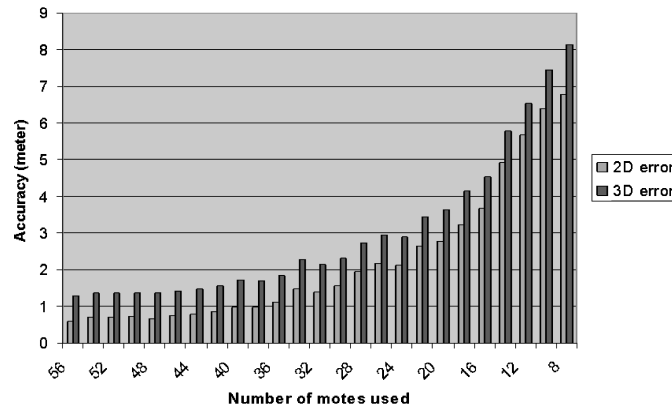


Fig. 13. Localization accuracy vs. number of sensors used.

Since N was decreased by two at a time and we stopped at eight nodes, we tested 250 different sensor network configurations.

We consider a shot undetected if there are less than six sensors detecting a muzzle blast. As the numbers of sensor decreased, so did the number of successfully detected shots as shown in Figure 12. Hence, the data in Figure 13 only uses the successfully detected shots. The diagram indicates that the error has an exponential characteristic. Close to our original setup, the error hardly increases. At 36 nodes, the average 3D error is still less than 2 meters. Beyond this point, however, the accuracy starts to decrease rapidly.

The raw results could lead to a premature conclusion that we could decrease the node density by 40% and still get very good accuracy. However, there are other considerations. Node failures decrease sensor density over time, so the planned deployment length needs to be considered. It is not enough to measure the acoustic events; the data also needs to be propagated back to the base station. There must be enough nodes to ensure a connected network with redundancy for robustness and good response time.

The overall accuracy of the system during the field tests in an urban environment indicates its tolerance to multipath effects. Of the 171 shots used in the analysis above, the average rate of bad measurements, that is, TOA data measured by individual sensor nodes that were not consistent with the final shooter location estimate, was 24%. In our experience, the vast majority of erroneous TOA data were due to multipath.

It is possible to solve the TDOA-based localization problem analytically, for example, as in Mahajan and Walworth [2001], where the constraints from measurements are converted to a linear equation system. This solution requires five measurements to determine the 3D position of a source, but it is straightforward to extend the solution in Mahajan and Walworth [2001] for more sensor readings. The solution of the over-determined equation system provides a least-squares estimation of the shooter location. We used this approach to evaluate our sensor fusion technique.

To compare the accuracy of the fusion algorithm to that of the analytical solution, field sensor measurements of 46 shots with known positions were used

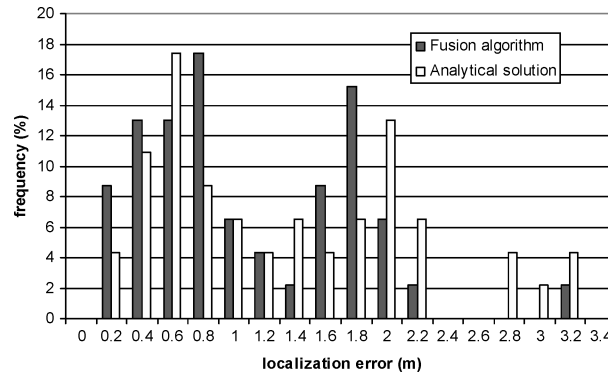


Fig. 14. Histogram of the localization errors using the fusion algorithm and the analytical solution.

as test cases. In the first test all the bad measurements resulting from multipath effects or sensor failure were removed from the data set. Each of the remaining set of good measurements was consistent with the known shooter position with time error less than 0.5 ms for each sensor reading. The shooter positions were estimated using both methods. The accuracies of the two solutions were very close to each other, as the histogram of errors shows in Figure 14. The mean 3D localization error for the fusion algorithm and the analytical solution were 1.2 m and 1.3 m, respectively, for the 46-shot test set. The difference is much less than the sensor position and the reference shooter position measurement errors, thus the performance of the two solutions can be considered to be equally good in this artificial test scenario.

In the previous test, the input contained only correct measurements. In practical cases, however, inconsistent measurements are present primarily due to multipath effects, even after careful pre-filtering of the sensor readings. To illustrate the sensitivity of the methods to measurement errors, bad sensor readings were added back to the input data set from the previously removed bad data set. For each shot, 2^B test sets were generated by combining the good measurement set with all possible combinations of the bad sensor readings containing B measurements. The number of good and bad sensor readings varied between 8 and 29, and between 1 and 10, respectively. Using all the 46 shots, 325 experiments were generated as test cases. Figure 15 shows the performances of the two methods, as a function of the ratio of the bad and good measurements.

It is clearly visible that the precision of the analytical solution was severely degraded when bad measurements were present, even in relatively small ratio. The fusion algorithm, however, was able to successfully eliminate the bad measurements, and its performance was the same as in the first test, independently of the ratio of the bad and good measurements.

6.4 Shockwave Fusion Results

The shockwave fusion algorithm was tested in the U.S. Army Aberdeen Test Center, in December 2004. The sensor network of 60 motes covered an 80×80 m area. The motes were placed on surveyed positions with an estimated

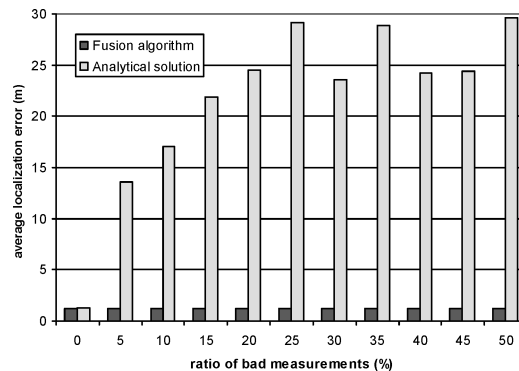


Fig. 15. The average localization error vs. the ratio of bad and good measurements.

accuracy of 20 cm. About half of the motes were on the ground, the rest were placed at elevations up to 6 m. Shooter positions were also surveyed. They were located up to 100 m from the edge of the sensor field. Two locations were in the basket of a cherry picker at approximately 12 m height. Various targets were placed inside and outside the sensor field opposite from the shooter positions. In one particular experiment, 12 shots were fired over the middle of the network shooting approximately 100 meters from the edge of the sensor field, so there were sensors on each side of the trajectory. The average azimuth error was 0.66 degrees, the average elevation error was 0.61 degrees, and the average range error was 2.56 meters. Another 11 shots were fired from the same distance near the edge of the network, so there were no or only a few sensors on one side of the trajectory. The average error increased to 1.41 degrees in azimuth, to 1.11 degrees in elevation and to 6.04 meters in range.

The latency of the shockwave-based fusion algorithm is somewhat greater than that of the muzzle-blast-based technique: the calculation of a single-shot trajectory takes about 3–4 seconds on a 3 GHz PC.

Multiple simultaneous shots have also been tested with mixed results. The typical test involved two shots only. About half the time, the system correctly localized both trajectories. There were cases, however, when three trajectories were found. Two of these were typically very close to each other and one of the true trajectories. This error can happen when there is more than double the number of detections needed for localizing a single trajectory. In such a case, the error value corresponding to a subset of the detections may be smaller than the one involving all the detections for one trajectory. The error function and the genetic algorithm need to be adjusted to avoid this situation.

Finally, the problem we refer to as “trajectory inversion” was also observed numerous times. Consider the two trajectories shown in Figure 16. One side of the generated shockwaves is almost the same. If we have sensors such that only this side is being measured, the wrong trajectory can be identified. In fact, the only way to differentiate the two possible sources of the half shockwave front is to measure its curvature caused by the deceleration of the projectile. However, that would mandate highly accurate sensor placement and very precise time synchronization. This can only be achieved by a

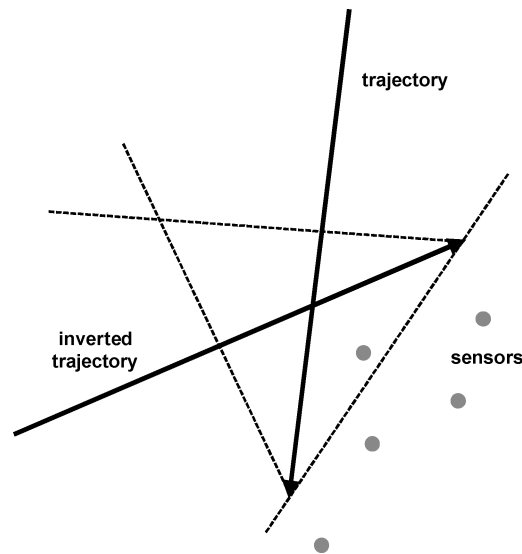


Fig. 16. Trajectory inversion problem.

microphone array with fixed and known microphone separation and a shared clock.

The area of correctly localizing multiple concurrent long range shots based on shockwave measurements needs further research. Unfortunately, we only had the opportunity to conduct a single field test. Our project ended in December 2004 and we have not worked on improving the presented techniques since.

The performance of the system is comparable to the published results of commercial centralized countersniper systems. BBN's *Bullet Ears* system utilizes one or two small arrays of microphones, providing estimates of the caliber, speed and trajectory of the projectile, and also a range estimate for the shooter. The average accuracy of the azimuth and elevation estimators is approximately 1.2 and 3 degrees, respectively, while the distance estimator's accuracy is approximately 1.6% [Duckworth et al. 1996]. The similar *Pilar* system uses two microphone arrays achieving bearing and range accuracy of $\pm 2^\circ$ and $\pm 10\%$, respectively [Pilar Systems].

6.5 Shockwave Fusion Sensitivity

To test the sensitivity of the fusion algorithm to the dominant sensor location error, we have conducted a preliminary experiment using the 12 shots fired in the middle of the network described above. The measured TOA results were used with modified sensor locations, similarly to the experiment in Section 6.3. The results for elevation accuracy look very similar to those of the azimuth summarized in Figure 17.

For each maximum additional sensor location error, each shot was tested ten times using different sensor locations. Therefore, each bar in the figure represents 120 experiments. The apparent insensitivity of the fusion to moderate sensor location errors is encouraging for the future development of the system.

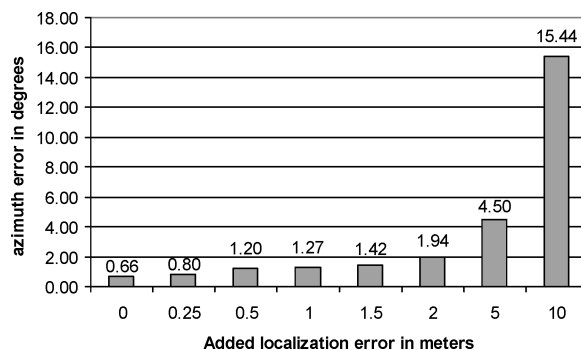


Fig. 17. Azimuth error vs. sensor localization error.

7. CONCLUSIONS

The article presented a comprehensive description of a real world application of wireless sensor networks. The success of the system is demonstrated by the fact that it has been transitioned to the industry. The fact that the system outperforms existing traditional systems in many respects in an extremely challenging application domain is a proof that WSN technology can indeed fulfill its great promise.

ACKNOWLEDGMENTS

The authors would like to thank Vijay Raghavan, Keith Holcomb, Al Sciarretta, David King, Tony Mason, Béla Fehér, James Davis, Tamás Lédeczi, Janos Sztiapanovits, Nancy Graves and Ben Abbott for their valuable contribution to this work. This work would not have been possible without the help and dedication of the people at the U.S. Army Dismounted Battlespace Battle Lab at Ft. Benning and the U.S. Army Aberdeen Test Center. We are grateful to Jim Reich and David Culler for their very constructive comments on earlier papers on the same topic.

REFERENCES

- BALOGH, G., LÉDECZI, A., MARÓTI, M., AND SIMON, G. 2005. Time of arrival data fusion for source localization. In *Proceedings of The WICON Workshop on Information Fusion and Dissemination in Wireless Sensor Networks (SensorFusion 2005)*, Budapest, Hungary, July.
- BLICKLE, T. AND THIELE, L. 1995. A comparison of selection schemes used in genetic algorithms. Tech. Rep. TIK-Report 11, Swiss Federal Institute of Technology (ETH), Zürich, Switzerland, May.
- CHEN, J. C., HUDSON, R. E., AND YAO, K. 2002. Maximum-likelihood source localization and unknown sensor location estimation for wideband signals in the near-field. *IEEE Trans. Sig. Proc.* 50, 1843–1854.
- CHEN, J., YIP, L., ELSON, J., WANG, H., MANIEZZO, D., HUDSON, R., YAO, K., AND ESTRIN, D. 2003. Coherent acoustic array processing and localization on wireless sensor networks. *Proc. IEEE*. 91, 1154–1162.
- DUCKWORTH, G. L., GILBERT, D. C., AND BARGER, J. E. 1996. Acoustic counter-sniper system. In *Proceedings of SPIE International Symposium on Enabling Technologies for Law Enforcement and Security*.

- ELSON, J., GIROD, L., AND ESTRIN, D. 2002. Fine-grained network time synchronization using reference broadcasts. *ACM SIGOPS Oper. Syst. Rev.* 36, SI, 147–163.
- GANERIWAL, S., KUMAR, R., AND SRIVASTAVA, M. B. 2003. Timing-sync protocol for sensor networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems (SenSys, 2003)* (Los Angeles, CA), 138–149.
- HILL, J. AND CULLER, D. 2002. Mica: A wireless platform for deeply embedded networks. *IEEE Micro*, 22, 12–24.
- HUANG, Y., BENESTY, J., AND ELKO, G. W. 2000. Passive acoustic source localization for video camera steering. In *Proceedings of the IEEE ICASSP 2000*. Turkey, June. IEEE Computer Society Press, Los Alamitos, CA, 909–912.
- KUSY, B., DUTTA, P., LEWIS, P., MARÓTI, M., LÉDECZI, Á., AND CULLER, D. 2005. Elapsed time on arrival: A simple and versatile primitive for canonical time synchronization services. *Int. J. Ad Hoc Ubiq. Comput.* In press.
- LÉDECZI, A., VÖLGYESI, P., MARÓTI, M., SIMON, G., BALOGH, G., NADAS, A., KUSY, B., AND DORA, S. 2005. Multiple simultaneous acoustic source localization in urban terrain. In *Proceedings of the 4th International Conference on Information Processing in Sensor Networks (IPSN 2005)* (Los Angeles, CA, Apr.).
- MAHAJAN, A. AND WALWORTH, M. 2001. 3-D position sensing using the differences in the time-of-flights from a wave source to various receivers. *IEEE Trans. Robotics Automat.* 17, 91–94.
- MARÓTI, M. 2004. The directed flood routing framework. In *Proceedings of the 5th ACM/IFIP/USENIX International Conference on Middleware* (Toronto, Ont., Canada, Oct.). ACM, New York, 99–114.
- MARÓTI, M., KUSY, B., BALOGH, G., VÖLGYESI, P., NADAS, A., MOLNAR, K., DORA, S., AND LÉDECZI, A. 2005. Radio interferometric geolocation. Accepted at *ACM Sensys 2005* (San Diego, CA, Nov.).
- MARÓTI, M., KUSY, B., SIMON, G., AND LÉDECZI, A. 2004. The flooding time synchronization protocol. In *Proceedings of the 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys 2004)* (Baltimore, MD, Nov.). ACM, New York, 39, 4–9.
- MOROZ, S. A., PIERSON, R. B., ERTER, M. C., BURCHICK SR., D. A., AND IPPOLITO, T. 1999. Airborne deployment of and recent improvements to the viper counter sniper system. In *Proceedings of the IRIS Passive Sensors* (Feb.).
- PILAR SYSTEMS. <http://www.army-technology.com/contractors/surveillance/metravib/>.
- SALLAI, J., BALOGH, G., MARÓTI, M., LÉDECZI, A., AND KUSY, B. 2004. Acoustic ranging in resource-constrained sensor networks. In *Proceedings of ICWN '04* (Las Vegas, Nv, June).
- SIMON, G., MARÓTI, M., LÉDECZI, A., BALOGH, G., KUSY, B., NADAS, A., PAP, G., SALLAI, J., AND FRAMPTON, K. 2004. Sensor network-based countersniper system. In *Proceedings of the 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys 2004)* (Baltimore, MD, Nov.). ACM, New York, 1–12.
- STOUGHTON, R. 1997. Measurement of small caliber ballistic shock waves in air. *JASA* 102 2, Pt. 1.
- VICK, A., STILLION, J., FRELINGER, D. R., KVITKY, J. S., LAMBETH, B. S., MAROUI, J., AND WAXMAN, M. C. 2000. Aerospace operations in urban environments: Exploring new concepts. In *RAND MR-1187*.
- XBOW. <http://www.xbow.com/>.
- XU, S. AND YANG, X. 2002. A review on interval computation—Software and applications. *Int. J. Comput. Numer. Anal. Appl.* 1, 149–162.

Received July 2005; accepted August 2005