



## Dynamic Networking and Smart Sensing Enable Next-Generation Landmines

*William M. Merrill, Lewis Girod, Brian Schiffer, Dustin McIntire, Guillaume Rava, Katayoun Sohrabi, Fredric Newberg, Jeremy Elson, and William Kaiser*

### EDITOR'S INTRODUCTION

Tragically, current-generation landmines linger indefinitely, threatening devastating injury for decades after hostilities cease. William Merrill and his colleagues discuss how to improve landmines for military and civilians using dynamic networking, networked sensors, and intelligent command and control. Their new-generation system can dynamically reconfigure itself to better protect military forces during warfare and also be deactivated and safely removed post bellum, preventing needless tragedies to civilians in ensuing years.

The pervasive networking and self-management capabilities shown in this department have broad applicability to environmental monitoring, such as the nearshore networks Rob Holman and his colleagues described in the October–December 2003 issue. Similarly, we can easily imagine geological monitoring of active sites or forest fire management using dynamically configuring devices by airdropping distributed networks into areas that are difficult, dangerous, or expensive to reach by land.

—Vince Stanford

To address the hazards to civilians, recent landmine systems have incorporated command detonation by remote operator and self-destruct capabilities (see the “Landmines through the Ages” sidebar). However, these additions didn’t let the mines adaptively counter breach attempts and might have actually increased the need to monitor mine location to allow command detonation. Nor did they allow opening breaches for desired safe passage. To provide these enhanced military capabilities and better protect civilians, future landmine systems should include

- Networked communication so mines can collaborate and communicate field status to an operator and enable

remote control as dictated by policy.

- Embedded processing for local decisions, such as what constitutes a breach, sensor fusions indicating disappearance of neighboring mines, and actions in response to changes in perceived status.
- Dynamically configurable communication to locate neighboring mines, and distributed collaborative algorithms to operate without external infrastructure (such as power or central communication) and to maintain gracefully degrading operational capability as the field is destroyed.

To improve military performance and prevent civilian casualties, the US military is integrating wireless networking and autonomous sensing into next-

generation landmines. These new systems will increase target discernment, enhance battle-shaping capabilities, and facilitate disposal or self-destruction of next-generation minefields when they’re no longer needed. From June 1999 to April 2002, our team developed wireless sensor network technology to build and test an example of this kind of system in the DARPA Applied Technology Office’s Self Healing Minefield program. Each SHM node builds and maintains awareness of its neighbor’s location and state so that a field of nodes can block, divert, or allow safe passage of specific vehicles.

### SELF-HEALING MINEFIELD REQUIREMENTS

A minefield, particularly in antivehicle roles, must block, divert, or trap vehicles. It presents a statistical obstacle in that its effectiveness is determined by the collective action of all mines, along all collective approaches. A way to improve system effectiveness is for the mines to collaborate and ensure their statistical response. Additionally, you might limit soldier oversight for minefield replacement. So, as much as policy allows, the minefield should autonomously configure and maintain its blocking, diverting, or trapping capability. Also, you should have communication within the minefield to enable user control when desired, such as in opening lanes for safe passage of

## LANDMINES THROUGH THE AGES

Greek historian Thucydides reported in *The Peloponnesian War* that primitive landmines were deployed as early as the fourth century B.C. These consisted of large clay pots buried in roads and fields, with a shallow earth covering, where enemy cavalry was expected to maneuver. While not explosive, they were still quite dangerous, as the camouflaged pots would collapse under a charging horse's weight, which almost invariably led to a broken leg for the horse and a thrown rider. Because terracotta can remain intact for millennia while buried, some of these primitive landmines might remain hazardous even today over 2,000 years later.

Explosive landmines are reported as early as the 15th century and were in wider use in the American Civil War in the 19th century. The modern mechanically fused antipersonnel mine first appeared in World War I, originating in Germany, but soon used by all belligerents. With the invention of the tank to break the stalemate of trench war, the anti-tank mine appeared as a countermeasure. To accomplish their purpose, antitank mines had to be large, but enemy infantry could then easily detect and move them. This prompted the deployment of large numbers of antipersonnel mines (hundreds of millions by all sides in World War II) to protect the antitank mines. These were typically deployed in defensive situations because they were individually buried by hand and not easy to redeploy dynamically.

After World War II, the technology advanced with new forms and new deployment methods such as artillery-fired or air-dropped cluster munitions. Armies could very quickly deploy these in offensive operations in great numbers and could insert them to separate cooperating units, deny paths of retreat, or maneuver in an enemy's rear areas. The

US first used this type in the Vietnam conflict, after which the technique was again widely copied and later used by the Soviet Union in the Afghanistan conflict. That era's landmines left no practical way to document the minefields for later removal and their low metal content made them hard to detect. Before her untimely death, Diana, Princess of Wales, devoted substantial portions of her public life to drawing public attention to the problem of these lingering and hard to detect landmines (see [www.landmines.org.uk](http://www.landmines.org.uk)).

In February 2004, US Assistant Secretary of State for Political-Military Affairs Lincoln Bloomfield announced a new Bush Administration policy on landmines recognizing the problem that previous generation devices have caused. This policy aims to commit the US to "eliminate persistent landmines of all types from its arsenal" and "a substantial increase in funding for our humanitarian mine action programs worldwide." He estimated that there are approximately 60 million persistent landmines remaining in the ground from various conflicts worldwide.

The policy has four major thrusts:

- All persistent mines will be eliminated by 2010, with destruction of such devices in the arsenal to begin within two years.
- Metallic content of all mines will be high enough for current demining equipment to detect.
- A concerted effort is underway to develop nonpersistent alternatives incorporating self-destructing, self-deactivating technologies and controls.
- It commits to a 50 percent increase in funding for humanitarian mine action programs starting in 2005 (see [www.state.gov](http://www.state.gov)).

friendly vehicles, and it provides redundancy in verification of mine removal or self-destruction.

SHM program requirements include a system-wide vehicular obstacle scalable to thousands of nodes. Each node must act to maintain this obstacle, even as large numbers of nodes are destroyed. So SHM nodes require capabilities for networking to collaborate with their neighbors, monitor changed neighbor status at 10-second intervals, determine the location of their neighbors to meter level accuracy, and move to reconstitute the obstacle if it's breached. In addition, nodes must function without an external infrastructure, such as GPS for location, external communication sites for coordination, or continuous oversight. They must also fit current antitank minefields deployment mechanisms (7.6 cm high, 12 cm diameter cylinder) and operate over

periods of days. Because the program is ongoing, nodes also have to provide a development environment that would let our research team most effectively meet evolving goals.

The SHM program offered general goals, such as identifying breaches in the field with low false-alarm rate, but no design guidance on how to achieve them. So the system had to be robust to often-intermittent wireless links and reliably differentiate the disappearance of nodes during a breach from lost packets due to radio frequency (RF) fading. Our system met all SHM goals as demonstrated at the program completion in April 2003.

Begun in 2000, our research aimed to create a development platform and a fielded system demonstrating a self-healing network and autonomous distributed processing. We developed an embedded computing platform, a soft-

ware infrastructure to enable development and testing of large collections of nodes, and a suite of software services to meet the application needs. We subcontracted to Science Applications International Corp. (SAIC), which managed the program and developed high-level mine control algorithms. Other team members included the Universal Propulsion Company, which developed the rocket engines used to enable mines to move and reconstitute a robust obstacle, and Ensign Bickford, which developed the safing and firing (SAF) system that controls each node rocket engine. Here, we focus on embedded networked computing platforms for a distributed system that Sensoria Corporation developed.

### SHM HARDWARE

Each SHM node contains a central processing unit; interfaces to subsys-

## APPLICATIONS

tems; an RF communication system; multichannel acoustic, magnetic, and acceleration sensing; and eight miniature rocket engines (see Figure 1). Unlike the microcontroller-based mote sensor nodes used for wireless sensors in many systems, our SHM node architecture provides significant embedded processing capability and a flexible development environment, which was critical to developing and deploying a multi-layer distributed system. Also, our SHM system provides a truly peer-to-peer architecture, with each node providing the same capabilities as every other node, and operates independently of external infrastructure. Table 1 details the hardware and software components.



**Figure 1. Self-healing minefield node shows wear from a system demonstration including hopping.**

## SHM SOFTWARE

On powering up, each SHM node determines if it can find nearby nodes, sets up a communication network to allow collaboration with identified neighbors, determines relative position of the 100 or so nearest neighbors, monitors their status via the network, and moves in response to each neighbor's disappearance to maintain a robust obstacle. To enable this operation, we at Sensoria developed the suite of software described in Table 2. Each software component contributes to that goal. The system represents a collaborative, layered application, with each software component building on those below. The software architecture is

**TABLE 1**  
Summary of SHM hardware features incorporated into each SHM node.

Component	Description	Specifications
Processor	Hitachi SH4 7751	300-MIPS CPU, 1.1Gflop floating-point unit.
Processor memory	16-Mbyte Flash, 64-Mbyte RAM	Available on the OS in various partitions.
Operating system	Linux	2.4.16 kernel extended with the framework for user space devices. Ported to the Hitachi SH4.
Preprocessor	Power control, subsystem interface	Subsystem monitoring and power control. Enables processor-to-subsystem communications.
Integrated power conditioning	Supplies power to each subsystem	0.15 W to 5.1 W adjustable power consumption using rechargeable, 7.4 V, 3.5 AH Lilon batteries.
Integrated acoustic subsystem	Analog-to-digital system supporting four separate acoustic input (microphone) and output (speaker) channels.	Two dual AC97 CODECs: <ul style="list-style-type: none"> <li>• Two stereo sigma delta 48-Ks/s inputs, with 46-dB programmable gain.</li> <li>• Two stereo 18-bit DAC outputs with 95-dB programmable attenuation.</li> </ul>
Digital I/O	14 lines of general-purpose I/O (GPIO) and four LEDs.	Subsystem control and development support.
Integrated antennas	Four 2.4-GHz Nitinol quarter-wave monopole antennas.	<ul style="list-style-type: none"> <li>• 2.4–2.49 GHz.</li> <li>• Switches between two on top and two on the bottom</li> </ul>
Mobility unit	Eight miniature solid propellant rocket engines. Four on top and four on the bottom.	Each engine provided a 10-m hop with a 2-m high trajectory. Hop directionality was determined by firing one of four engines, 90 degrees apart.
Package	Case and internal mounting enabling the nodes to survive multiple hops.	12-cm diameter, 7.6-cm high cylinder.
Orientation sensor	Three-axis magnetometer and accelerometers.	Detected node orientation and movement due to tampering or unintended node mobility.
Development access	Detachable 10/100-Mbps Ethernet daughter card and serial console.	Developer interface and monitoring.
Dual radios	Two 2.4-GHz ISM band frequency-hopping spread spectrum (FHSS) radio modems (Cirronet WIT2410).	Dual radios enabled scalable mesh network topology for automation.
Integrated antennas	Four Nitinol whip antennas, two on top and two on the bottom	Independent antennas for each radio switched so those on the topside are always used.

TABLE 2  
Software modules operating on the SHM communication units.

Software subsystem	Description
Distributed networking services	This subsystem forms and maintains the mesh network. Components include network assembly, link monitoring, a multihop distributed database, and point-to-point and point-to-multipoint route creation and monitoring, as well as efficient message flooding. <sup>2</sup>
Time synchronization	Radio heartbeat packets used to provide microsecond timing synchronization between nodes over multiple network hops. <sup>3</sup>
Relative position determination	Each node determines the relative position of the nearest 100 nodes. From the local range and angle data shared over the SHM network, positioning is separated into an acoustic ranging to determine range and angles between local nodes, and multilateration to enable every node to build a large set of node coordinates. <sup>4</sup> The SHM multilateration and ranging algorithms operating on the SH4 processor utilized up to 100 percent of the available processing resources.
Breach healing	Each node determines its own response to neighbor node disappearances when notified by the up-and-down detection software.
GUI interaction and debug support	Software modules coordinate with multiple software GUIs and offer extensive debugging and logging capabilities. Specific interfaces provide the state of each safing and firing (SAF) system and report what thrusters will fire over the radio, provide system operation specifics via the Ethernet development interface, and enable reporting and node control via radio or Ethernet interfaces.
Up and down detection	This software coordinates with the orientation and distributed network services to verify and report node disappearances. Operators monitor node status solely via the network, and disappearances are detected with corroborated loss of wireless links. The software distinguishes between a system-wide network failure and the loss of specific nodes within a breach.
Hardware module drivers	Interfaces to each hardware subsystem include the radios, the SAF to control firing the rockets (as indicated by healing), the orientation subsystem, the preprocessor power control, and the acoustic-sampling section.
Flexible development system operating environment	We built each SHM node on the Linux OS with additional layers added to streamline device and process interaction, enhance watchdog functionality, monitor and control process operation tailored to an energy constrained autonomous system, and provide multiple monitoring points within each software module. We developed these within a soft-state approach to enable robust operation and graceful degradation in the presence of faults.

designed using open system interfaces, which provide multiple interface layers over the Linux kernel. This lets the SHM system software components work together to support the application goals as a system, even if individual nodes might appear to be ineffective when considered locally.

### SHM PERFORMANCE

We designed our SHM system as part of an autonomous infrastructure-free weapon system. It integrated a combination of autonomous acoustic geolocation, self-assembling reliable wireless ad hoc communication, limited mobility, distributed sensing, and autonomous distributed collaboration. While many researchers are developing and demonstrating systems in each of these areas,<sup>1,5,6</sup> the SHM system is unique in its level of integration and the robustness

needed for a self-contained, autonomous, and wirelessly networked weapon system. Particularly, it offers multiple levels of graceful degradation of the network formation and routing, acoustic geolocation, and statistical response algorithms that support the system as a whole.

The SHM system operates autonomously outdoors once deployed at the specified density of one node per 70 m<sup>2</sup>. When powered up, the nodes create a local area network within 2 to 5 minutes, determine the relative positions of up to 100 neighbors in about 15 minutes, and collaboratively monitor other nodes' state and operation for at least eight hours. Each SHM node autonomously joins the local networks, then uses the created multihop topology to maintain time synchronization, coordinate acoustic geolocation, share acousti-

cally determined range and angle information between nodes to enable relative geolocation, and share status information on all of their neighbors.

### Testing

We demonstrated the SHM's operation over a month of testing at Fort Leonard Wood, Missouri, from early March to early April 2003. During these tests, from 50 to 95 individual SHM nodes operated using internal batteries for multiple trials per day, demonstrating and monitoring the autonomous networking, relative position determination, collaborative decision-making, and identification of node faults. The nodes were set up outdoors, in a field 35-m deep by up to 190-m across (for 95 nodes). We obtained the results of system operation presented here through noninvasive Ethernet connections to a

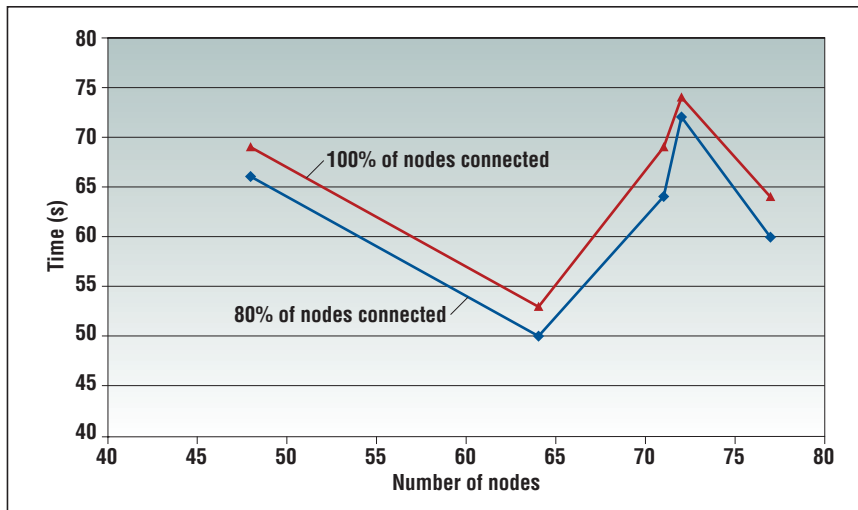


Figure 2. Times for the network to initially stabilize. We recorded only one test at each number of nodes.

fraction of the nodes, reporting data on autonomous system operation. To demonstrate the SHM system's operation, we'll discuss a few of the results demonstrated at Fort Leonard Wood.

When powered up, each of the two radio modems integrated into each SHM node form independent local RF clusters, with a base modem coordinating up to nine remote modems. Because each node participates in two independent clusters and clusters are formed to overlap, the SHM nodes can autonomously form a scalable interconnected network. When the SHM nodes first power up, they autonomously discover their wireless neighbors and collaborate to build a robust network topology and routes.

Figure 2 illustrates the time for the initial SHM network topology to stabilize to include connections between 80 and 100 percent of the nodes in selected tests. This figure shows the first stable connections formed by at least one radio on each node. Additional link-by-link refinement occurs until the network stabilizes on both radios—generally by 300 s. The lack of correlation of test size and network formation time in this figure results from the parallel formation of the SHM dual-cluster

topology and from the scalability of the developed network assembly algorithm. The network formation time variation between tests of the same size at Fort Leonard Wood (within 20 s), resulting from randomized timer values in the formation algorithm, exceeded the variation for increasing numbers of nodes past a size of 50.

To enable a one-way time-of-flight acoustic ranging variation from 3 to 175 ms depending on node separation, and inter-microphone, the SHM system leverages a reference broadcast time-synchronization mechanism.<sup>3</sup> We used this time synchronization on each receiver time-stamping incoming broadcast packets and then comparing their time stamps. The system assumes that each time stamp varies with Gaussian error about an average (after outliers are filtered out) and records the variance of the remaining time stamps as the root mean square synchronization error to each node that receives some subset of the same broadcasts. We monitored these RMS error values to a precision of one microsecond.

#### Node synchronization

Figure 3 shows an example of the distribution of RMS error for four tests

between all directly synchronized nodes. To synchronize nodes that do not receive the same physical layer broadcast messages, local SHM node time synchronization between remotes coordinated by the same base is shared across multiple RF hops. This approach enables acoustic ranging between nodes, independent of the local RF communication topology.

To autonomously locate its nearest 100 neighbors, each SHM node uses acoustic time-of-flight measurements. While similar in concept to some other systems,<sup>7</sup> this system is a fully distributed peer-to-peer system with every node providing the same capability. Within the SHM system, the range and angle between pairs of nodes are collected, then distributed within 10 RF hops (about 100 nodes). The system uses one-way time-of-flight measurements between nodes to calculate range, with the arrival time differences between each of four microphones on each node serving to calculate relative angles.

#### Accuracy

Example accuracy demonstrated 85 percent of range measurements within 1 m of ground truth and 70 percent of angle measurements within  $\pm 20^\circ$  of ground truth. Once it obtains ranges and angles, the system refines the initial guess at location these provide on each node via a least-squares optimization to determine node locations within two dimensions.<sup>4</sup>

Figure 4 shows one node's perspective on the error seen in this geolocation process. The figure uses blue circles to display node 235's location table for multiple runs at Fort Leonard Wood over a few days. Red triangles in this figure represent those nodes' ground truth locations. The figure shows a few reported outlier node locations (outside the expected  $170 \times 35$ -m box in which nodes were deployed), that result from initial error in node placement; they are corrected as the nodes acquire more ranges. The figure is made up of multiple coordinate location snapshots, at

10-minute intervals, during multiple tests. Because the multilateration improves over time with more ranges, early coordinate locations from the first snapshot are less accurate than later ones. The figure demonstrates another feature of the SHM geolocation: the multilateration algorithm's tendency to rotate the node positions around the node at the origin.

We found significant variation in the time to collaboratively determine coordinates for a field of SHM nodes. Tests with more than 50 nodes increased this time resulting from the increased time to complete multilateration on the integrated Hitachi SH4 7751 32-bit microprocessor. High time variability to determine a location table resulted from high winds, temperature inversion, and high noise conditions. Generally, a field of 50 nodes required from 10 to 20 minutes for all nodes to determine and share their positions. With up to 100 nodes, this time variability increased from 10 to 30 minutes. In all tests run, all nodes acquired location tables because even in the worst conditions, the nodes' operation was robust enough to collect data, with an exponential time back off, until all nodes could calculate their neighbors' positions.

### Reducing complexity

The SHM network supports various distributed services, including geolocation and collaborative status monitoring. Looking at the level of total SHM traffic, the average radio traffic sent over each link—averaged over all links between base and remote radios—was 139 bytes/s. To place this RF traffic in perspective, the system's maximum sending capacity is 5.27 Kbytes/s for bases and 3.23 Kbytes/s for remotes.

Figures 2 and 4 provide a sampling of the SHM system's performance. We developed this system over three years to meet application goals that, while well defined, required a significant research effort. Within that development, a number of features were critical to the system's creation. Given the

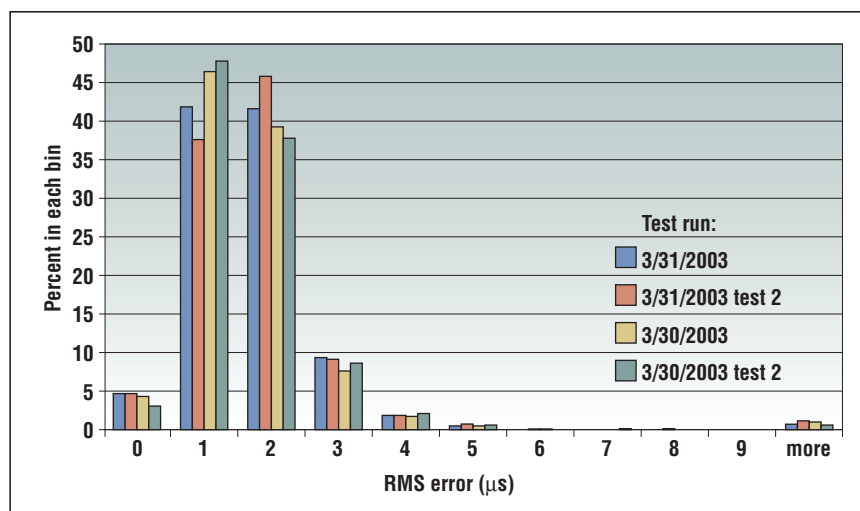


Figure 3. CPU-to-CPU RMS synchronization error reported for four sets of synchronized nodes.

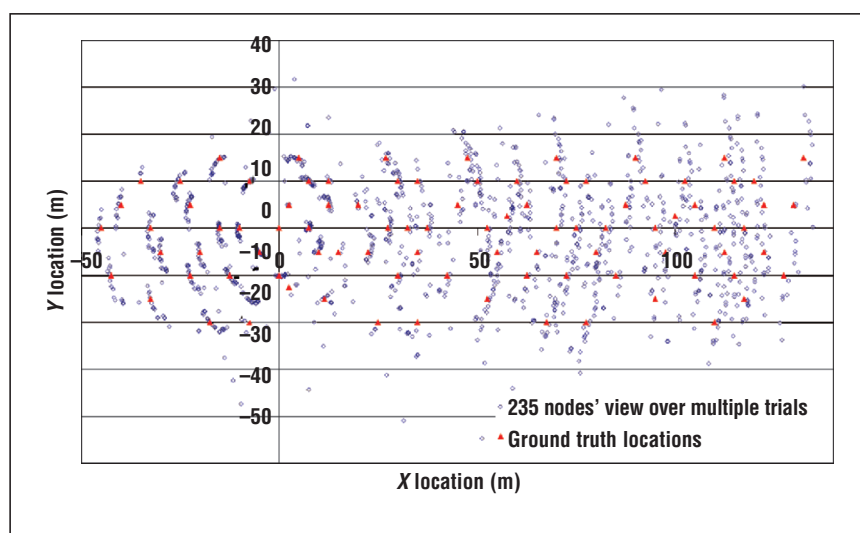


Figure 4. Multiple location tables collected by node 235 over a week (in blue) compared with hand-measured node locations (in red).

competing requirements for an autonomous system that involved collaboratively forming and monitoring large numbers of nodes over an intermittent wireless network, the layered component solutions combined to create a complex system. This complexity resulted from the interaction of the networking, geolocation, status monitoring, and distributed response algorithms. Even though we intentionally kept each algorithm straightforward

to facilitate implementation and testing, the interaction of layers of algorithms produced a large degree of system complexity.

The required system-level robustness also dictated that each system component provide a graceful degradation path and layers of error mitigation. Various system features—from the hardware preprocessor that monitored the main processor to detect system lockups and restart the nodes, to the implementation

## APPLICATIONS

of error detection and resilience at each step of the multilayer geolocation solution—incorporated redundancy and resilience to component failures and errors. Within the SHM development, we saw a critical need in autonomous system development for these two features: designing in the node flexibility and processing capability to support a complex solution that will change during development, and requiring multiple layers of error detection and resilience in each component solution.

The SHM system shows how a wireless networking and autonomous use of sensing can augment and influence a weapon system's scope. This system demonstrates an application that can be supported with state-of-the-art embedded processing, energy-constrained communication, and distributed software components. We believe the SHM system illustrates one of the first examples of an applied wireless networked embedded system not dependent on an external infrastructure and whose application space requires significant distributed computing and communication capability.

To integrate the SHM wireless sensor network we've described into a deployable weapon system requires significant additional engineering integration to ensure a reliable, easily deployable system that would operate in all desired deployment environments. This future work would include

- Further integration of power control and duty cycling to extend the system's lifetime to multiple days
- An integrated warhead (this development system did not include a warhead) and firing algorithms
- Enhanced communication security and the use of packet radios operating in an approved military frequency band and waveform
- Additional component integration to further reduce size and meet required US military environmental susceptibility and mechanical standards

Currently, with the completion of the DARPA-funded effort, this work is not going forward directly to create an SHM weapon system. Rather, some of the demonstrated SHM technologies are under consideration for integration in and support of other weapon systems. However, this article provides an example of the capability of wireless networking and embedded computing to transform staid, well-known systems, such as a field of antivehicle landmines. ■

## ACKNOWLEDGMENTS

This work was supported by DARPA/ATO contract DAAE30-00-C-1055. We also acknowledge the substantial guidance provided by Thomas Altschuler currently with Rockwell Scientific, Glenn Rolader with SAIC, and Kent Carson of the Institute for Defense Analysis.

## REFERENCES

1. R. Szweczyk et al., "Habitat Monitoring with Sensor Networks," *Comm. ACM*, vol. 47, no. 6, June 2004, pp. 34–40.
2. K. Sohrabi et al., "Methods for Scalable Self-Assembly of Ad Hoc Wireless Sensor Networks," *IEEE Trans. Mobile Computing*, to be published in 2004.
3. J. Elson, L. Girod, and D. Estrin, "Fine-Grained Network Time Synchronization Using Reference Broadcasts," *Proc. Fifth Symp. Operating Systems Design and Implementation (OSDI 2002)*, USENIX, 2002, pp. 147–163.
4. W. Merrill et al., "Autonomous Position Location in Distributed, Embedded, Wireless Systems," *Proc. IEEE CAS Workshop Wireless Communications and Networking*, IEEE Press, 2002, pp. 1–8.
5. I.F. Akyildiz et al., "Wireless Sensor Networks: A Survey," *Computer Networks*, vol. 38, no. 4, 2002, pp. 393–422.
6. G.J. Pottie and W.J. Kaiser, "Wireless Integrated Network Sensors," *Comm. ACM*, vol. 43, no. 5, 2000, pp. 51–58.
7. N.B. Priyantha et al., "The Cricket Compass for Context-Aware Mobile Applications," *Proc. 7th Int'l Conf. Mobile Computing and Networking*, ACM Press, 2001, pp. 1–14.

**William W. Merrill** is the lead wireless architect at Sensoria Corp. Contact him at Sensoria Corp., 15950 Bernardo Center Dr., Ste. J, San Diego, CA 92127; williamm@sensoria.com.

**Lewis Girod** is a PhD candidate at UCLA and works part time at Sensoria Corp. as a senior development engineer. Contact him at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; lew@sensoria.com.

**Brian Schiffer** is a senior design engineer at Sensoria working on hardware for embedded systems. Contact him at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; brschiff@sensoria.com.

**Dustin McIntire** is a graduate student at UCLA working on his MS in electrical engineering and a senior design engineer at Sensoria. Contact him at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; dustin@sensoria.com.

**Guillaume Rava** is senior software engineer at Sensoria. Contact him at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; grava@sensoria.com.

**Katayoun Sohrabi** is the lead network architect at Sensoria Corp. Contact her at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; sohrabi@sensoria.com.

**Fredric Newberg** is the lead hardware architect at Sensoria Corp. Contact him at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; fredric@sensoria.com.

**Jeremy Elson** is a postdoctoral researcher at UCLA's Center for Embedded Networked Sensing. Contact him at the UCLA Center for Embedded Networked Sensors, Los Angeles, CA 90095; jelson@cs.ucla.edu.

**William Kaiser** is a professor in the Department of Electrical Engineering in the Henry Samueli School of Engineering and Applied Sciences at UCLA and cofounder of Sensoria Corp. Contact him at Sensoria Corp., 200 Corporate Pointe, Ste. 100, Culver City, CA 90230; kaiser@sensoria.com.