

Sensor Networks for Emergency Response: Challenges and Opportunities

A common protocol and software framework could integrate devices such as wearable vital sign sensors, handheld computers, and location-tracking tags into disaster response scenarios. CodeBlue, a new architecture, allows wireless monitoring and tracking of patients and first responders.

Sensor networks, a new class of devices, have the potential to revolutionize the capture, processing, and communication of critical data for use by first responders. Sensor networks consist of small, low-power, and low-cost devices with limited computational and wireless communication capabilities. They represent the next step in wireless communication's miniaturization, and their power and size make it feasible to embed them

into wearable vital sign monitors, location-tracking tags in buildings, and first responder uniform gear.

Sensor nodes' extreme resource limitations represent new challenges in protocol design, application development, and security models. We developed CodeBlue, a common software infrastructure, to address these challenges. Code-

Blue integrates sensor nodes and other wireless devices into a disaster response setting and provides facilities for ad hoc network formation, resource naming and discovery, security, and in-network aggregation of sensor-produced data. We designed CodeBlue for rapidly changing, critical care environments. To test it, we developed two wireless vital sign monitors and a PDA-based

triage application for first responders.

Additionally, we developed MoteTrack, a robust radio frequency (RF)-based localization system, which lets rescuers determine their location within a building and track patients. Although much of our work on CodeBlue is preliminary, our initial experience with medical care sensor networks raised many exciting opportunities and challenges.

The CodeBlue infrastructure

Integrating a range of wireless devices with varying capabilities into medical, disaster response, and emergency care scenarios (see the related sidebar) raises new challenges for the devices' interoperability. Sensor nodes' extremely limited communication and computational capabilities exacerbate these challenges. Some specific issues that arise include

- *Discovery and naming.* Establishing communication pathways between vital sign sensors and receiving devices—for example, a PDA that an emergency medical technician (EMT) carries—requires flexible device discovery. Device naming should also be application centric. Rather than using low-level network addresses, a device might want to request information on all vital sign sensors of a certain type or only on those devices emitting data in a certain range.

Konrad Lorincz, David J. Malan, Thaddeus R.F. Fulford-Jones, Alan Nawoj, Antony Clavel, Victor Shnayder, Geoffrey Mainland, and Matt Welsh
Harvard University

Steve Moulton
Boston University

Emergency Scenarios and Sensor Networks

Knowing how to both monitor and deal with a large number of casualties is key to disaster response scenarios. If first responders can't rapidly triage the injured and severely injured in a coordinated manner, the large numbers could quickly overwhelm emergency field personnel and hospital staff and prevent them from providing quality trauma care.¹ Quickly identifying and stratifying the most severely injured patients poses unique challenges, as does efficiently monitoring and transporting victims.

The triage process has traditionally occurred at the hospital gate (typically at the entrance to the Emergency Department). There, an emergency physician or experienced trauma care surgeon stratifies patients with one or more triage tools based on mechanism of injury, physiologic criteria, injury site and severity, preexisting disease, age, and survival expectation. Overtriage, a common problem, ties up valuable resources that would otherwise go to more severely injured patients. Emergency personnel must decide as early as possible which patients will benefit most from transport to a dedicated trauma center and which patients require less immediate attention.

We envision sensor network nodes playing a variety of disaster response roles:

- In-field patient triage and tracking
- Temporary storage of individual patient information, including on-scene physical exam findings, treatment types, and treatment response
- Simultaneous physical environment monitoring
- Tracking first responders' and patients' location and status

The most compelling application is wireless vital sign monitoring of multiple victims at a disaster scene. First responders would place

wireless, low-power vital sign sensors on each patient. These sensors would relay continuous data to nearby paramedics and emergency medical technicians, who would use mobile computers or PC-based systems (for example, in ambulances) to capture additional patient data. Field personnel could therefore monitor and care for several patients at once, yet still be alert to sudden changes in any particular patient's physiologic status.

Using *decision support* to guide trauma care is another sensor networking opportunity in disaster response settings. The quality of disaster-related trauma care is inversely proportional to its caseload. By continuously feeding patient information from the field into a centralized decision support system, responders could gain a global view of a mass casualty situation and establish a greater semblance of order. This would make triage in the field and at the hospital gates interactive, allowing better coordination of the out-of-hospital caseload with critical, hospital-based trauma facilities and resources.

In a December 1999 fire in Worcester, Massachusetts, six firefighters died after they became lost in a six-floor warehouse; four of the firefighters were attempting to locate the first two victims, who might have already died.² Equipping first responders with vital sign monitors would avoid placing additional rescuers at risk because they'd know if victims were beyond medical help. Additionally, location-tracking systems based on radio frequency, ultrasound, or another technology could aid rescuers in determining their own location as well as that of others.

REFERENCES

1. E.R. Frykberg and J.J. Tepas III, "Terrorist Bombings: Lessons Learned from Belfast to Beirut," *Annals of Surgery*, vol. 208, no. 5, 1988, pp. 569–576.
2. A. Flint and S. Milligan, "'Your Tragedy Is Ours': President Joins Thousands in Saluting Fallen Firefighters," *Boston Globe*, 10 Dec. 1999.

You must decentralize the discovery process to avoid any single point of failure; it would be inappropriate to rely on a central directory server.

- **Robust routing.** Disaster scenario devices might need to communicate with other devices outside their immediate radio range. For example, you might have nodes distributed over a large area and rescuers frequently moving. Ad hoc routing techniques extend the effective communication range by having devices relay messages for one another. Notably, communication patterns are typically multicast, so a given vital sign sensor might need to report its data to multiple receiving nodes.

- **Prioritization of critical data.** Communications bandwidth is extremely limited on low-power radios such as those based on the IEEE 802.15.4 and ZigBee standards. When many devices share this bandwidth, you must give critical data—such as vital signs from an arresting patient or an SOS message from a trapped firefighter—priority over other traffic.
- **Security.** Security is a major concern for wireless communication systems in general, and even more so in disaster response. A dynamically changing population of patients and rescuers at a disaster site requires efficient establishment of security credentials. In addition, the security architecture can't

assume a predeployed public key infrastructure (PKI) or that all devices have sufficient computational power to run expensive cryptographic protocols.

- **Tracking device locations.** Location awareness is an important aspect of disaster response: incident command operations depend heavily on tracking rescuer and victim locations. Given the many wireless devices in a disaster site, you should be able to use GPS, RF signals, ultrasound, or some other technique to track patient and rescuer device locations.

CodeBlue (see Figure 1) comprises a suite of protocols and services that lets

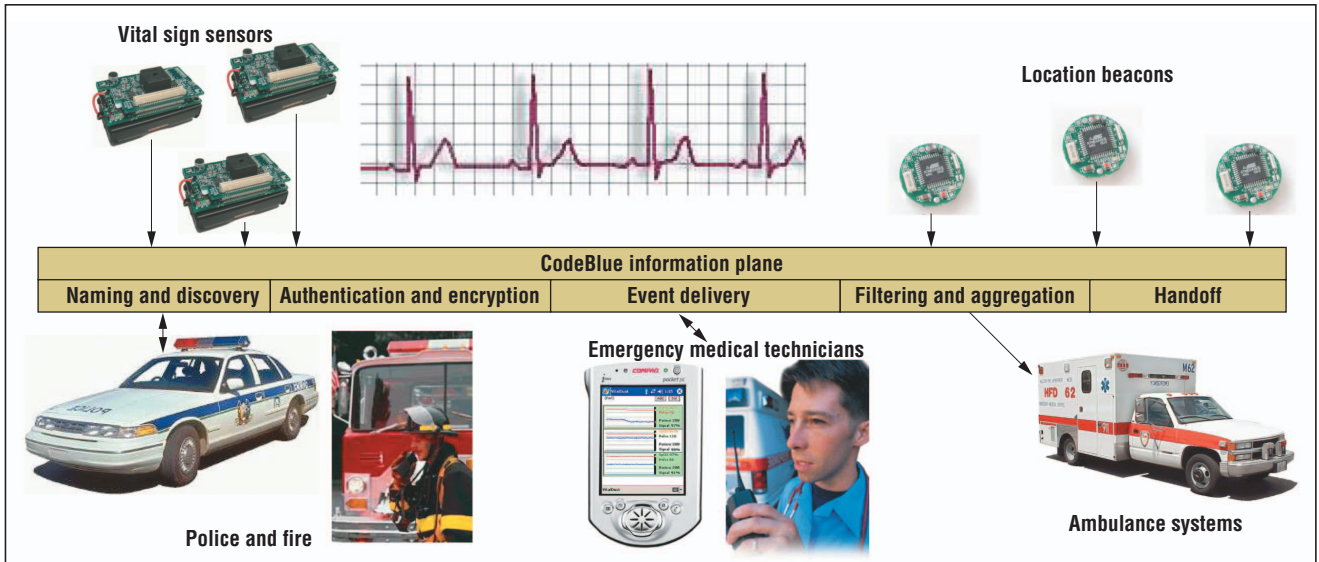


Figure 1. The CodeBlue infrastructure.

many types of devices (wireless sensors, location beacons, handheld computers, laptops, and so forth) coordinate their activities. We think of CodeBlue as an “information plane,” which lets these devices discover each other, report events, and establish communication channels. CodeBlue incorporates a flexible naming scheme; robust publish and subscribe routing framework; authentication and encryption provisions; and services for credential establishment and handoff, location tracking, and in-network filtering and aggregation.

The pervasive systems community has worked toward the seamless integration of many computational devices for some time, but integrating low-power, low-capability wireless sensors into emergency and disaster response demands new approaches. For example, unlike traditional middleware architectures, CodeBlue must run on sensor network devices with extremely limited resources. Traditional approaches based on Remote Procedure Call (RPC), Java virtual machines, or mobile agents are generally inappropriate for this domain.

Several projects share similar goals to CodeBlue's. The Patient Centric Network project (<http://nms.lcs.mit.edu/projects/pcn>) is developing a common

architecture for sensors in hospitals, such as those found in operating rooms. This work isn't specifically focused on low-power, wireless sensors in a disaster response environment. Several projects, including Agent Based Casualty Care (www.cs.dartmouth.edu/~afiske/abccare), are developing wearable physiological sensors. Many of these projects focus on a particular application and sensor suite, rather than a more general framework for wireless medical sensing. Additionally, most of these systems use PDA-class systems with 802.11b, which have very different power and bandwidth characteristics than the low-power sensor nodes that we envision.

We've completed some promising early designs of several CodeBlue components.

Wireless vital sign monitors

Wireless sensor networks are typified by small, low-power, and low-capability devices, often called *nodes*, a term referring to a speck of dust. The Mica2 mote, developed at UC Berkeley,¹ is one of the most popular sensor node designs. It's based on a 7.3-MHz Atmel ATmega128L embedded controller with 4 Kbytes of RAM and 128 Kbytes of ROM. The Mica2 includes a low-power, single-chip

radio—the Chipcon CC1000—capable of operating at 76.8 kbps with a practical indoor range of approximately 20–30 meters. The Mica2 measures 5.7 cm × 3.2 cm × 2.2 cm and uses two AA batteries that will last for up to a week if the device is powered continuously. However, you can extend its lifetime to months or years through careful duty cycling. The Mica2 runs a specialized operating system, called TinyOS, that addresses the sensor nodes' concurrency and resource management.

These devices represent a very different design point from other mobile technologies, such as PDAs, which have far greater CPU speeds, memory sizes, and communication capabilities. Sensor nodes' limited bandwidth and computational power precludes the use of common Internet protocols and services such as the TCP/IP, DNS, and Address Resolution Protocol (ARP). Additionally, very different technology trends drive these devices' size, cost, and power consumption than those in the handheld or wearable computer space. Sensor nodes already exist that integrate all of Mica2's capabilities onto a single 5 mm² chip.²

Wireless sensing and communication have the potential for broad applications in medicine. Today, it's possible to obtain heart rate, oxygen saturation,

end-tidal CO₂, and serum chemistries measurements, including serum glucose, with small, noninvasive sensors. We expect that, over time, an increasing array of sensors with sophisticated capabilities will become available. Companies such as Nonin and Numed have wireless vital sign sensors based on Bluetooth technology, while Radianse developed an RF-based location-tracking system for hospital use. Other research projects include the European Commission's wide-ranging Mobi-Health Project, which provides continuous monitoring of patients outside the hospital environment by developing the concept of a 3G (third-generation)-enabled "Body-Area Network." The potential applications will save lives, create valuable medical research data, and cut medical services costs.

To demonstrate wireless sensor nodes' use in disaster response, we've developed two mote-based vital sign monitors (see Figure 2): a pulse oximeter and a two-lead electrocardiogram (EKG) monitor.³ The pulse oximeter captures a patient's heart rate and blood oxygen saturation (SpO₂) by measuring the amount of light transmitted through a noninvasive sensor attached to the patient's finger. EMTs use these standard vital signs to determine a patient's general circulatory and respiratory status, which are among the first vital signs taken. We based our mote-based pulse oximeter on Smith-BCI's commercially available daughterboard that attaches to the Mica2 mote, which transmits the heart rate and SpO₂ data periodically (about once a second).

Our mote-based EKG continually monitors the heart's electrical activity in more severely injured patients. For example, a patient with internal bleeding might require cardiac monitoring to determine that the heart rate and rhythm are within acceptable limits. You can use EKG signals to detect arrhythmia (abnormal heartbeat rhythm) or ischemia (lack

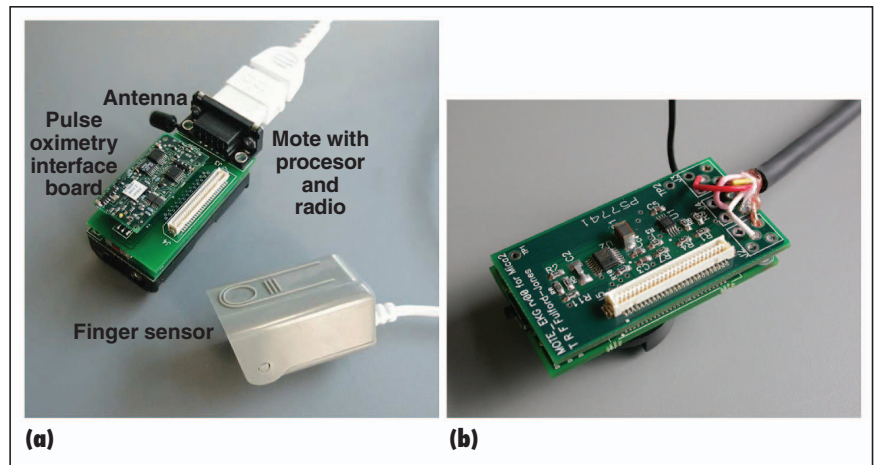


Figure 2. Mote-based (a) pulse oximeter and (b) two-lead electrocardiogram.

of blood flow and oxygen to the heart), both of which point to potentially serious conditions. A custom-built circuit board attached to the Mica2 mote captures a trace of the heart's electrical signals through a set of leads attached to a patient's chest. The circuit board captures information at a rate of 120 Hz, compresses it using a differential encoding scheme, and transmits it via the mote's radio.

Handheld computers carried by first responders can receive and visualize multiple patients' vital signs. Figure 3 shows our PDA-based triage application, which displays real-time data from multiple patients and can report a combination of audible and visible alerts should a patient's vital signs exceed a predetermined range. Additionally, the real-time data collected by our PDA-based application can be passed on to a patient care record application such as 10Blade's iRevive. iRevive provides a PDA-based patient care record that EMTs can use for recording patient history, identification, and other observations, as well as any interventions such as intubation, medications, or fluid resuscitation. We're currently working with 10Blade to integrate our two PDA-based applications.

Security implications

Security is an important factor to wireless sensor networks' success and accep-

tance in medical and disaster response applications because patients' medical records must remain private. In hospitals and other clinical settings, medical devices must ensure the privacy of patients' medical data in accordance with the Health Insurance Portability and Accountability Act of 1996. For in-



Figure 3. PDA-based multiple-patient triage application. The screen shows real-time vital sign (heart rate and blood oxygen saturation) data from three patients.

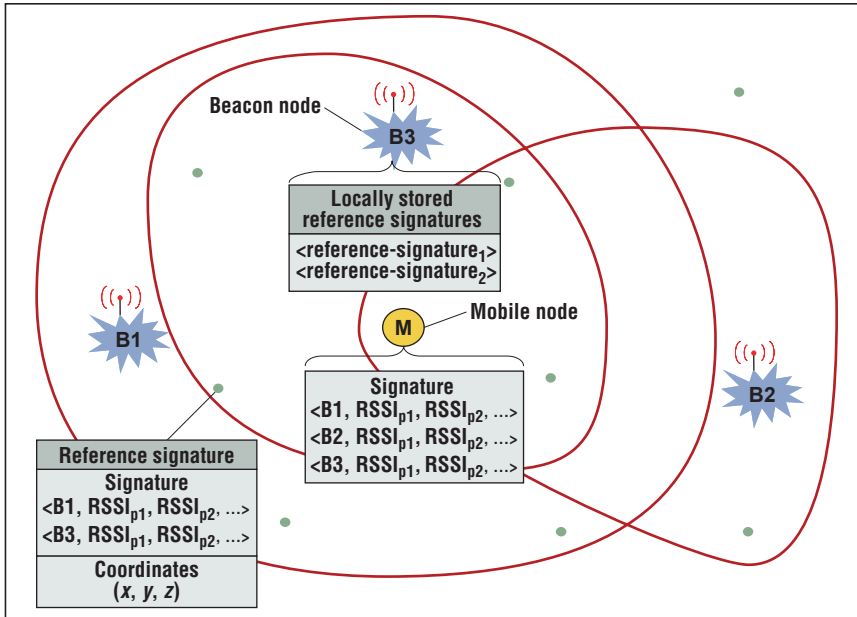


Figure 4. The MoteTrack location system. B1, B2, and B3 are beacon nodes, which broadcast beacon messages at various transmission powers (p_1 , p_2 , and so forth.). Each beacon node stores a subset of all reference signatures. M is a mobile node that can hear from all three beacon nodes. It aggregates beacon messages received over some time period into a signature. The circular areas marked by red perimeters indicate the areas in which messages can be received from the corresponding beacon node.

field use, devices must also defend against adversaries bent on capturing, spoofing, or inducing denial-of-service against the system. A large disaster scenario will involve rescue workers and ambulances from many organizations. It's unreasonable to assume that all organizations have exchanged security or configuration information (keys, certificates, and so forth) ahead of time. It's also impractical to require that rescue personnel arriving on the scene spend time typing in passwords, logging into databases, or any typical authentication methods used for conventional computing. Therefore, we need an architecture that supports an ad hoc network security model that doesn't require manual configuration and that self-organizes based on the set of participating devices. The system must cope with nodes joining and leaving the system—for example, as EMTs arrive on the scene and patients are transported off-site to a hospital. The security model must also provide seamless credential handoff, allowing one first responder to give access rights to another, without relying on preexisting relationships between the two.

In traditional distributed security, you authenticate principals by contacting a trusted authority that's responsible for maintaining up-to-date information on what each principal's access rights are. However, in a disaster response scenario, it might be impossible to access an outside authority due to lack of preexisting infrastructure or inability to contact off-site systems. In life-threatening situations, it's never acceptable to deny a legitimate user data that might save someone's life. In such situations, a best-effort security model might be appropriate, making strong guarantees when external authorities can be contacted and making weaker guarantees during periods of poor connectivity or infrastructure loss.

Although public key cryptography can address many of these problems, employing such an approach can introduce several technical difficulties. Sensor nodes' limited resources are ill-suited for most straightforward implementations of public key cryptography. For example, with only 4 Kbytes of memory, the Mica2 can store no more than a few 1,024-bit RSA keys.

We've been exploring Elliptic Curve Cryptography as an alternate public key cryptography scheme.⁴ ECC uses smaller keys and is more computationally efficient than RSA. An ECC key size of 163 bits is computationally equivalent to a 768-bit RSA key, and we can implement ECC using only integer arithmetic, which is much more feasible on low-power microcontrollers without hardware floating-point support. Our implementation of ECC on the Mica2 can generate a key in 35 seconds, which, while far from negligible, is still acceptable if we perform key generation infrequently. You can use this approach to generate symmetric keys for symmetric-key encryption schemes such as TinySec,⁵ thereby amortizing the overhead over many transmissions.

In the future, we'd like to explore taking advantage of the diversity of device capabilities in a disaster setting. It might be possible to develop schemes in which the PDAs and laptops on the scene perform expensive security computations. This approach doesn't completely address the problem, however, because sensor nodes will still need to determine which devices to trust for offloading these computations.

RF-based location tracking

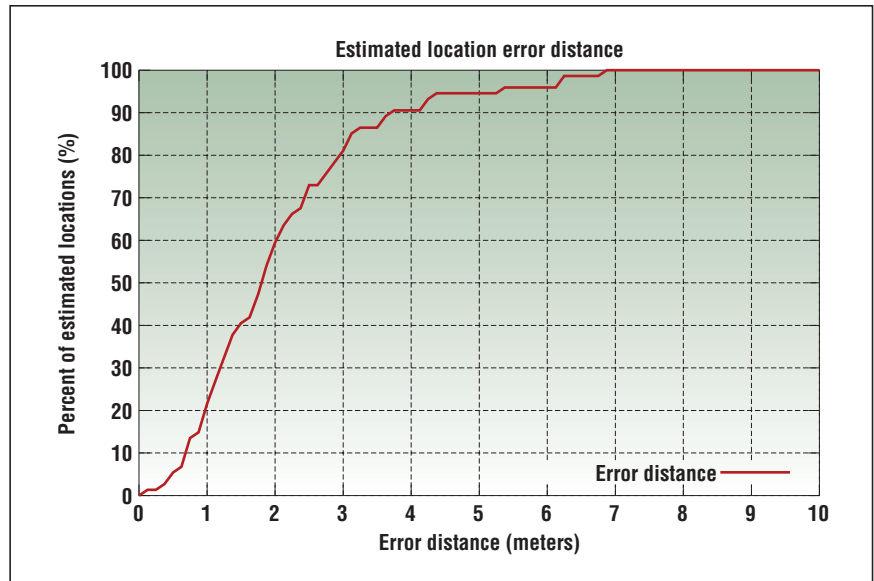
Tracking patients and rescue personnel is another important wireless tech-

Figure 5. MoteTrack location estimation accuracy. Eighty percent of the location estimates are within 3 m of the true location. The data is for 74 location estimates collected over one floor of the Computer Science building at Harvard University, measuring roughly 1742 m² (18751 ft²). We distributed a total of 282 reference signatures to 20 beacon nodes.

nologies application in disaster response. For example, in a mass casualty incident, you can place tractable vital sign sensors on many patients. This lets you quickly locate a patient who suddenly requires immediate attention—an essential part of a successful triage. Firefighters entering a large building often cannot see because of heavy smoke coverage and have no a priori notion of building layout. By installing wireless, battery-operated RF beacons in a building in advance of a fire, firefighters and rescuers could use a heads-up display to track their location and monitor safe exit routes.⁶ Likewise, an incident commander could track multiple rescuers' locations in the building from the command post. Such capabilities would have greatly improved FDNY rescue operations on 11 September 2001, according to the McKinsey reports.⁷

We're developing an RF-based location tracking system, called MoteTrack, specifically designed for disaster response.⁸ It operates using the low-power, single-chip radio transceivers found in sensor network nodes, which rescue personnel can easily wear or embed in wearable vital sign sensors. MoteTrack operates in an entirely decentralized, robust fashion, providing good location accuracy despite partial failures of the location-tracking infrastructure.

With MoteTrack, you populate a building or other area with a number of *beacon nodes*, which can operate off battery power or use main power with



a battery backup. These beacon nodes can replace existing smoke detectors and serve as both wireless smoke detectors and location trackers. If the nodes track location infrequently, such as only in an emergency, and with careful duty cycling, the operation lifetime of beacon nodes running off batteries will resemble that of a battery-operated smoke detector. MoteTrack doesn't require additional hardware beyond the sensor node's radio and microprocessor. Beacon nodes, represented as stars in Figure 4, broadcast periodic beacon messages that consist of a tuple of the format {sourceID, powerLevel}. sourceID is the unique identifier of the beacon node, and powerLevel is the transmission power level used to broadcast the message.

Each mobile node that wants to use MoteTrack to determine its location listens for some period of time to acquire a signature. A signature consists of the beacon messages received over some time interval along with the received signal strength indication (RSSI) for each transmission power level. Finally, we define a reference signature as a signature combined with a known 3D loca-

tion (x, y, z).

MoteTrack uses a two-phase process to estimate locations: an offline collection of reference signatures followed by online location estimation. Once you've installed the beacon nodes, you use a mobile node to acquire a reference signature set at known, fixed locations throughout the building (shown as green dots in Figure 4). Afterwards, a mobile node can obtain a signature and send it to the beacon node from which it received the strongest RSSI to estimate its location. This approach resembles 802.11-based location-tracking systems, such as RADAR.⁹ However, unlike RADAR, MoteTrack is completely decentralized—that is, it runs entirely on small, low-power sensor nodes and doesn't require a back-end database to store reference signatures or perform location calculations. MoteTrack carefully replicates its reference signature set across beacon nodes such that each beacon node stores only a subset of all reference signatures. The beacon nodes themselves perform all data storage and computation using only the locally stored reference signatures. MoteTrack achieves a high level of robustness to beacon node

failure through its distributed architecture and by using an adaptive algorithm for estimating locations that is a function of the percent of locally failed beacon nodes. The high level of robustness to failure is an important factor for disaster response applications.

In an experiment consisting of 20 beacon nodes distributed throughout one floor of Harvard University's computer science building measuring 1742 m^2 (that is, a density of 0.011 (beacon nodes/ m^2)), MoteTrack achieved an 80th-percentile location accuracy of 3 meters over 74 separate location estimates and tolerated a failure of up to 40 percent of the beacon nodes with negligible increase in error. Figure 5 shows the distribution of location-tracking error. This accuracy is roughly equivalent to that of research and commercial 802.11-based location-tracking systems (www.radianse.com/products.htm),⁹ but doesn't require a powered infrastructure or connection to a network to track locations, is entirely decentralized, and is robust to failure. Ultrasound-based systems, such as Cricket,¹⁰ yield a higher degree of accuracy but need very dense beacon placement and line-of-sight beacon exposure. This might require you to carefully orient the receivers and is impractical for rescue operations.

Wireless sensor networks have the potential for enormous impact on many aspects of disaster response and emergency care. However, the medical community is large, diverse, and sophisticated, and federating these systems will require addressing a range of technical, scientific, social, financial, and legal issues.

Moreover, a host of technical challenges remain for integrating these devices into disaster settings. Power, computational capabilities, and communication

bandwidth limitations demand new approaches to software design in this regime. We're currently developing CodeBlue, which will integrate device discovery, robust routing, traffic prioritization, security, and RF-based location tracking. Our initial prototypes of these services show promise, and we're planning several deployments in simulated and real clinical settings. ■

REFERENCES


1. J. Hill et al., "System Architecture Directions for Networked Sensors," *Proc. 9th Int'l Conf. Architectural Support for Programming Languages and Operating Systems (ASPLOS 2000)*, ACM Press, 2000, pp. 93–104.
2. J. Hill, *System Architecture for Wireless Sensor Networks*, doctoral thesis, Dept. Electrical Eng. and Computer Sciences, Univ. of California, Berkeley, 2003.
3. T.R.F. Fulford-Jones, G. Wei, and M. Welsh, "A Portable, Low-Power, Wireless Two-Lead EKG System," to be published in *Proc. 26th IEEE EMBS Ann. Int'l Conf.*, 2004.
4. *IEEE Standard Specifications for Public-Key Cryptography*, Microprocessor and Microcomputer Standards Committee, IEEE CS Press, Jan. 2000.
5. C. Karlof, N. Sastry, and D. Wagner, "TinySec: Link Layer Security for Tiny Devices," www.cs.berkeley.edu/~nks/tinysec.
6. G. Slack, "Smart Helmets Could Bring Firefighters Back Alive," *Forefront*, Eng. Public Affairs Office, Univ. of California, Berkeley, Fall 2003; www.coe.berkeley.edu/forefront/fall2003/helmet.html.
7. *Increasing FDNY's Preparedness*, City of New York, 2004, www.nyc.gov/html/fdny/html/mck_report/index.shtml.
8. K. Lorincz and M. Welsh, *A Robust, Decentralized Approach to RF-Based Location Tracking*, tech. report TR-19-04, Division of Eng. and Applied Sciences, Harvard Univ., 2004.
9. P. Bahl and V.N. Padmanabhan, "RADAR: An In-Building RF-Based User Location and Tracking System," *Proc. IEEE Conf. Computer Comm. (INFOCOM 2000)*, IEEE CS Press, 2000, pp. 775–784.
10. N.B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The Cricket Location-Support System," *Proc. 6th Ann. ACM Int'l Conf. Mobile Computing and Networking (MobiCom 00)*, ACM Press, 2000, pp. 32–43.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.

the AUTHORS



Konrad Lorincz is a third-year PhD candidate in computer science at Harvard University. His research interests include distributed systems, networks, wireless sensor networks, location tracking, and software engineering. He received his MS in computer science from Harvard University. Contact him at 238 Maxwell Dworkin, 33 Oxford St., Cambridge, MA 02138; konrad@eecs.harvard.edu; www.eecs.harvard.edu/~konrad.



David J. Malan is a third-year doctoral student in computer science at Harvard University. His research interests include cybersecurity and worms. He received his SM in computer science from Harvard University. Contact him at Maxwell Dworkin, 33 Oxford St., Cambridge, MA 02138 USA; malan@eecs.harvard.edu; www.eecs.harvard.edu/~malan.

the AUTHORS



Thaddeus R.F. Fulford-Jones is a staff researcher in electrical engineering and computer science at Harvard University. His research focuses on novel integrations of biomedical sensors and motes. He received his MS in electrical engineering from Harvard University. He is a member of the IEE. Contact him at 209 Maxwell Dworkin, 33 Oxford St., Cambridge, MA 02138 USA; fulford@fas.harvard.edu; www.people.fas.harvard.edu/~fulford.



Alan Nawoj is a senior software engineer for Deloitte Consulting and a hardware and software engineer for Spotlight Mobile. His research interests include mobile computing challenges and wireless sensor networks. He received his MS in computer science from Harvard University. Contact him at 63 Highland Ave., Apt. 5, Cambridge, MA 02139; nawoj@post.harvard.edu.



Antony Clavel received his MS in computer science from Harvard University. His research interests include distributed systems and artificial intelligence, focusing on sensor networks and autonomous multiagent systems. Contact him at 1572 Massachusetts Ave., Apt. 26, Cambridge, MA 02138; clavel@post.harvard.edu.



Victor Shnayder is a second-year graduate student at Harvard University. His research interests include large-scale distributed systems, sensor networks, and security. He received his BSE in computer science from Princeton University. Contact him at 238 Maxwell Dworkin, 33 Oxford St., Cambridge, MA 02138 USA; shnayder@eecs.harvard.edu; www.eecs.harvard.edu/~shnayder.



Geoffrey Mainland is a second-year PhD student at Harvard University. His research interests include programming languages, distributed systems, networks, and intelligent control. He received his AB in physics from Harvard College. He is a member of the ACM. Contact him at 199 Mount Auburn St. #3, Cambridge, MA 02138; mainland@eecs.harvard.edu; www.eecs.harvard.edu/~mainland.



Steve Moulton is an associate professor of surgery and pediatrics at Boston University School of Medicine. He is the chief of pediatric surgery and director of pediatric trauma at Boston Medical Center and the founder of 10Blade. His research focuses on developing software and sensors for optimal care of the trauma patient. He received his MD from the University of Washington. He is a fellow of the American College of Surgeons and the American Academy of Pediatrics, Surgical Section. Contact him at Dowling 2419, One Boston Medical Ctr. Pl., Boston, MA, 02118; smoulton@bu.edu.



Matt Welsh is an assistant professor of computer science at Harvard University. His research interests include operating system, network, and programming-language support for massive-scale distributed systems, including Internet services and sensor networks. He received his PhD in computer science from the University of California, Berkeley. He is a member of the ACM and the IEEE. Contact him at 233 Maxwell Dworkin, 33 Oxford St., Cambridge, MA 02138 USA; mdw@eecs.harvard.edu; www.eecs.harvard.edu/~mdw.