

DistressNet: A Wireless Ad Hoc and Sensor Network Architecture for Situation Management in Disaster Response

*Stephen M. George, Wei Zhou, Harshavardhan Chenji, MyoungGyu Won, Yong Oh Lee, Andria Pazarloglou, and Radu Stoleru, Texas A&M University
Prabir Barooah, University of Florida*

ABSTRACT

Situational awareness in a disaster is critical to effective response. Disaster responders require timely delivery of high volumes of accurate data to make correct decisions. To meet these needs, we present DistressNet, an ad hoc wireless architecture that supports disaster response with distributed collaborative sensing, topology-aware routing using a multichannel protocol, and accurate resource localization. Sensing suites use collaborative and distributed mechanisms to optimize data collection and minimize total energy use. Message delivery is aided by novel topology management, while congestion is minimized through the use of mediated multichannel radio protocols. Estimation techniques improve localization accuracy in difficult environments.

INTRODUCTION

A hurricane comes ashore near Galveston, Texas in 2008. Hundreds of thousands of people are displaced as high winds, torrential rains, and flooding force them from their homes. A few thousand stay behind. Buildings are submerged or swept away as public utilities fail. Roads are blocked by debris and flood waters. Survivors find themselves trapped with no water, electricity, or communications.

In the aftermath of the storm, the first 36 hours are critical. A command post is assembled in a large field to manage disaster response. Police, fire, and medical personnel converge on the scene from scores of jurisdictions. Assigned areas to search, they spread throughout the disaster to assess the situation, assist survivors, and locate casualties. A medical triage facility treats victims as a complex logistics infrastructure is assembled. Teams of specialists work feverishly to restore electricity and water. At the peak of the response effort, the total number of responders exceeds 65,000.

Management of resources in a disaster is difficult. Rescue personnel spread across wide swaths of the disaster area. Situational awareness is maintained using hand-written tables and charts. Updates arrive in the form of verbal messages via radio or courier. In this environment wireless ad hoc and sensor networks (WASNs) can significantly enhance situational awareness by improving and automating updates, monitoring and reacting to status changes, and extending data communications across the entire disaster. However, while WASNs have made significant contributions in surveillance [1], target tracking [2], and healthcare [3], they have not achieved broad application in disaster response. Several challenges make their integration into this field difficult.

First, disasters are infrequent, and the location, communications requirements, and sensing needs of the next disaster cannot be predicted. Existing WASN systems often fill niche needs, like specialized chemical detection, that are ill suited for general deployment. Other solutions, like building monitoring systems, require pre-disaster installation and maintenance. To succeed, WASNs supporting disaster response must be extensible, flexible, adaptive, and designed to leverage and incorporate emerging technologies.

A second challenge involves scale and standardization. Small disasters, like a localized flood, normally involve resources from a single jurisdiction and are organized using simple ad hoc command and control (C2). Larger disasters, like a hurricane, may incorporate thousands of resources from scores of jurisdictions. To manage complexity at this scale, responders often organize following Incident Command System (ICS) guidelines [4]. To operate in this difficult environment, WASNs must provide standard architectures and composable, scalable networks structured to conform to ICS standards.

Lack of communications infrastructure is a third challenge faced in larger disasters. Entire regions suffer from degraded communications,

and remaining capacity is exhausted by the demands of victims. Disaster responders must arrive with their own communications. WASNs designed for disaster response must provide standalone, robust, and pervasive communications sufficient to support the geographic coverage and mobility requirements of its users.

A final challenge involves the occasionally congested but normally dispersed and disconnected nature of disaster response activities. Responders are highly mobile and often scattered across long distances. Groups coalesce to solve manpower-intensive problems but quickly separate to continue search and rescue (SAR). Sensors are deployed densely to monitor unstable structures or scattered widely to track dispersion of gaseous chemicals. Since all elements periodically report status and location, WASNs operating in this environment must be prepared to communicate opportunistically as conditions permit, and tolerate high mobility, long disconnections, and infrequent updates.

As a solution to these challenges, we propose DistressNet, a WASN architecture for disaster response. Contributions and novelties of this work include:

- An extensible, scalable, heterogeneous network addressing the challenges of disaster response by providing area-wide situational awareness through large-scale integration of static and mobile elements with common protocols and a robust support infrastructure
- Novel network management and configuration strategies that handle arrival and departure of elements, minimize congestion through the use of multi-channel communications, maximize message delivery with opportunistic and delay-tolerant protocols, and offer guaranteed levels of service for safety-critical messages
- Distributed collaborative sensing and robust, adaptive localization and location-based services that provide accurate, detailed data that enhances decision support, situational awareness, and general C2 activities

DistressNet is being implemented using custom and commodity sensors, mobile and static gateways capable of cross-protocol routing, and a variety of servers providing network services, analysis, and decision support. This article gives an overview of the DistressNet architecture including principal network elements and key software components. The second section presents related work. The third and fourth sections detail the system architecture and software components, respectively. The final section offers conclusions.

RELATED WORK

ICS [4] is a set of guidelines for organizing disaster response. It provides standardized but flexible mechanisms to guide the formation of collaborative teams capable of cross-jurisdictional coordination. ICS also provides a framework of common processes supporting integration of resources from different organizations into cohesive teams.

Sensor networks have been deployed to a

variety of challenging environments. ExScal [1] and VigilNet [2] are two such systems. ExScal deployed over 1200 elements in an outdoor monitoring application with a static preplanned topology. VigilNet, with approximately 200 nodes, used a static but unplanned topology in a similar effort. Both projects used mostly homogeneous hardware. DistressNet requires greater scalability, more heterogeneity, and higher degrees of mobility. AlarmNet [3], deployed in an assisted living environment, offers a heterogeneous network of static and mobile elements deployed across a confined area. Unlike DistressNet, AlarmNet relies on fixed infrastructure covering bounded areas. Code Blue [5], another medical system, relies on fixed infrastructure in mostly indoor areas. DieselNet [6], a bus-based delay-tolerant network (DTN) with a high degree of mobility along predictable routes, relies on the presence of a large amount of 802.11 infrastructure. In contrast, DistressNet assumes availability of no external infrastructure along its varied unpredictable routes.

DISTRESSNET SYSTEM ARCHITECTURE

DistressNet is a complex multilayer multiprotocol architecture displaying high degrees of mobility and varied connectivity. Figure 1 offers insight into the scope of the problems presented by a disaster. The inset map demonstrates how teams with distinct missions operate at a distance from the C2 area. Periodic reports to the C2, containing the location and status of personnel and equipment, are required to maintain situational awareness. As shown in Fig. 1a, there is often no direct network path between a team and the C2 area. Voice communications are used, but channels may be congested due to the number of transmitters. Data communications are enabled through the use of DTN protocols. Figure 1b offers a different problem. Although teams are still isolated from the C2, the presence of three SAR teams, each with a separate audio sensing network, causes local congestion. Multichannel medium access control (MAC) protocols allow the teams to coexist without interference while data communications with C2 are enabled using DTN protocols. The C2 area in Fig. 1c faces a different set of problems. C2 elements manage data from hundreds of workers and dozens of sensing networks. High-bandwidth connections are critical, but the network must incorporate DTN protocols in order to communicate with outstations.

In spite of the apparent complexity of its requirements, DistressNet uses simple components to form integrated networks. The composition of these elements is guided by a set of core principles.

Extensibility through the use of common protocols: Disaster response is composed of many small specialized elements assembled to meet the requirements of the situation. DistressNet is similarly designed to enable incorporation of new elements and networks through the use of common protocols like 802.11, 802.15.4, and IPv6, and simple network applications.

In spite of the apparent complexity of its requirements, DistressNet uses simple components to form integrated networks. The composition of these elements is guided by a set of core principles.



Figure 1. DistressNet, superimposed on an image of the aftermath of Hurricane Ike, Galveston, Texas, 2008: a) utility workers restore electrical and water service (BodyNet/TeamNet); b) search and rescue teams (BodyNet/TeamNet) look for victims in damaged buildings using sophisticated audio sensing suites (SenseNet); c) command and control, medical triage, and logistics work under an AreaNet umbrella. (Photo courtesy of the National Weather Service.)

Cross-layer optimization of MAC and routing: Communications within DistressNet are complex. Some segments of the network may be congested while others are frequently disconnected. Management of these extremes requires robust cross-layer MAC optimizations that leverage available spectrum and transport mechanisms for optimal communication with every node. Complex routing schemes involving multiple frequencies and delay-tolerant segments must be created, tracked, and updated with enough frequency to guarantee communications.

Composable localization that manages uncertainty: Every node in DistressNet is localized. Localization is critical to the situational awareness required by C2 elements. DistressNet incorporates composable localization using various techniques to reduce uncertainty and provide accurate locations.

Data aggregation and cooperative sensing to minimize energy use and network congestion: Sensing and localization consume energy. The communications required for these activities are even more energy-intensive. Within DistressNet, networks collaborate in their sensing and local-

ization tasks to minimize individual energy consumption and maximize network lifespan. They distribute sensing duties where possible and cooperatively communicate complete answers instead of raw data.

PRINCIPLE COMPONENTS

At its most basic level, DistressNet is composed of sensors that know or can determine their location, name, and network address. These sensors, often elements of more complex systems, organize into networks to perform sensing and support tasks. Divided by function, these networks represent individual sensors and people, static and mobile sensing suites, multi-element teams, and large-scale distributed systems. They range from minimally capable tracking devices to multiprocessor, multichannel, multiprotocol routers. DistressNet integrates these networks as shown in Fig. 1 and provides a common software architecture to enable sensing, communications, and localization (Fig. 2).

BodyNet is a body-worn wireless sensor network designed to monitor the health and status of its host (Figs. 1a and 2a). A BodyNet also

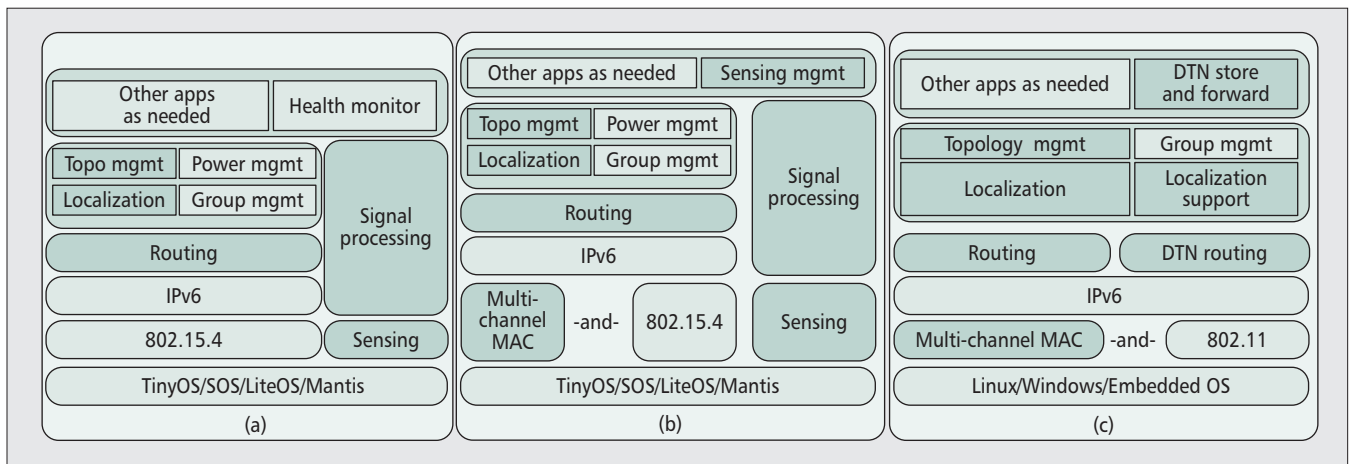


Figure 2. DistressNet software architecture (darker blocks represent contributions). a) BodyNet; b) SenseNet; c) VehicleNet and AreaNet.

provides localization for its hosts and sends updates to C2. First responders, support staff, and C2 elements within a disaster area are outfitted with BodyNets. Typically covering a range of meters using 802.15.4, a BodyNet may include additional sensors like specialized monitoring systems. Data from BodyNet nodes is transmitted through higher-level networks to C2 elements to support situational awareness applications.

TeamNet is a network composed of the BodyNets of the members of a single team, one or more VehicleNet nodes, and collocated equipment (Fig. 1a). A TeamNet normally represents a group performing a single role like SAR or medical support. Managed by a central node on the team’s vehicle, TeamNet operates in a full mesh to track and monitor team member status. TeamNet nodes provide each other with localization and communications support. Data within a TeamNet is consolidated and summarized before being forwarded to C2 elements.

VehicleNet is one of two backbone networks within DistressNet (Fig. 1a). Composed of nodes located on every vehicle deployed to the disaster, VehicleNet offers mobile routing and DTN to elements throughout DistressNet. VehicleNet nodes bridge 802.15.4 and 802.11 segments of DistressNet (Fig. 2c). In more densely populated areas VehicleNet nodes function as routers. In sparser areas they provide delay-tolerant relay services. Since every TeamNet includes at least one vehicle, mobile VehicleNet nodes provide the most common path between TeamNets and the network core.

AreaNet is the other backbone network within DistressNet (Fig. 1c). Composed of static nodes deployed to localized positions, AreaNet provides high-bandwidth 802.11-based mesh networking. Also capable of bridging 802.15.4 and 802.11 network segments, AreaNet has the capacity to carry large volumes of data and support bandwidth-intensive rescue operations involving imagery, audio, and video (Fig. 2c). AreaNet grows as responders take control of increasingly larger sectors of the disaster area.

SenseNet describes a variety of specialized networks within DistressNet (Fig. 1b). These

networks transmit sensed data and status updates to C2 elements via VehicleNet or AreaNet, and may be grossly divided into three categories: specialized sensing, monitoring, and tracking (Fig. 2b). Specialized sensing networks are employed in a variety of niche applications including audio monitoring, structural stability tracking, and distributed chemical sensing. Monitoring networks may be used to monitor the status of physical objects or people (e.g., victims in a triage area or water levels behind a levee). Tracking networks may be used to maintain visibility of items throughout the disaster area (i.e., specialized tools, radios, and generators).

DISTRESSNET SOFTWARE COMPONENTS

The integration of the elements discussed in the previous section and depicted in Fig. 1 form DistressNet. However, in order to sense, localize, and communicate, these elements rely on common software components. The software components shown with darker shading in Fig. 2 are described in detail in the next section.

SENSING

The primary purpose of any wireless sensor network is sensing. In DistressNet, collaborative distributed sensing provides local nodes and C2 elements with information required to maintain situational awareness. Sensing systems within DistressNet cover areas as small as two meters, in the case of a single human body, and as large as hundreds of meters, in the case of distributed chemical detection.

We propose a novel form of collaborative distributed sensing to enhance SAR. In order to detect trapped victims, simple wireless audio sensors are deployed around damaged buildings to cover areas inaccessible to rescuers. The sensors form a multihop mesh network and monitor their local audio spectrum. Upon detecting audio in the range of a typical human voice, 300 Hz–4 kHz, the sensors collaborate to record the sound. Since high-resolution sampling is not sustainable due to constraints on power consump-

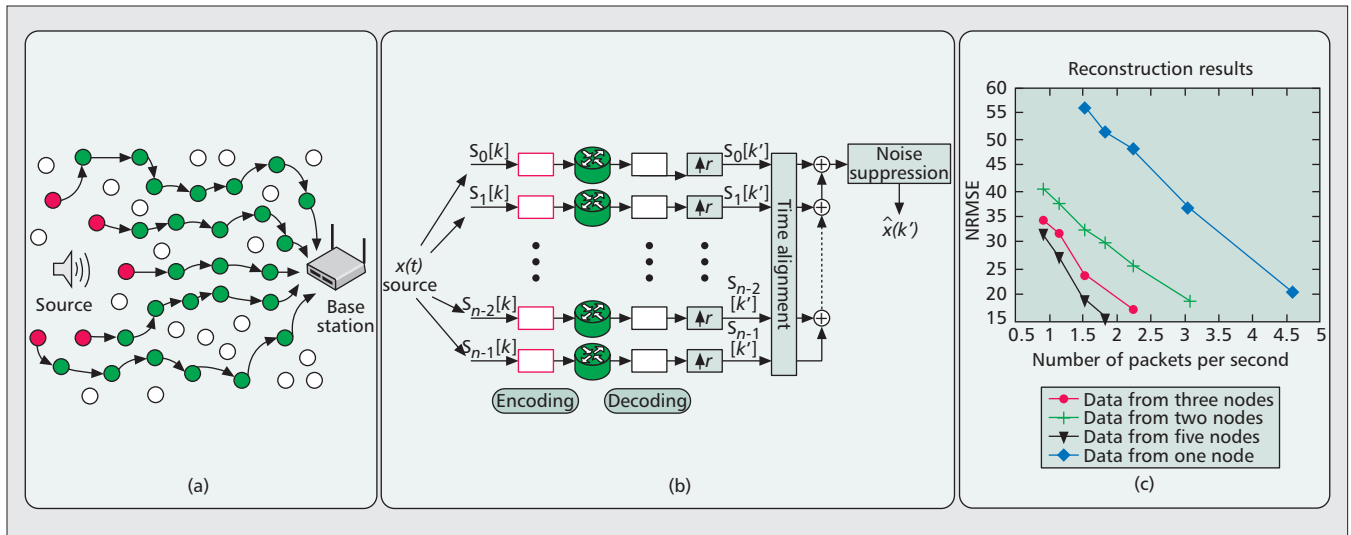


Figure 3. Distributed audio sensing: a) multiple sub-Nyquist streams carry sensed audio; b) sampling, transport, and reconstruction; c) audio reconstruction using various numbers of audio sensors and transmission rates.

tion and bandwidth, each sensor samples at a sub-Nyquist rate of 1 kHz and streams the collected audio over multiple routes to a central node for processing (Fig. 3a).

Figure 3b illustrates the sampling and reconstruction. The original audio source is denoted $x(t)$. Discrete-time signals $s_i(k)$, $i = 0, \dots, n - 1$ are noisy sub-Nyquist versions of $x(t)$ produced by each node after sampling (k denotes the time indices at sensor node sampling rate f_s and k' denotes the time indices for the higher-resolution signals at sampling rate $f_{s'} = f_s \cdot r$ after upsampling). To optimize $s_i(k)$ for transmission, 1-byte encoding provides a 2:1 compression ratio. For reconstruction, the received streams are decoded by the base station and upsampled to a common rate through interpolation to compensate for variable time delays. Time delay estimation using the resulting signals $s_n[k']$ aligns them properly on the time domain. The output signal is formed by superimposing the aligned signals and applying a denoising filter.

Analysis indicates that when multiple nodes cooperate to sample audio at low resolution, the results can approximate audio sampled by a single sensor at much higher resolution. Figure 3c uses the normalized root mean square error (NRMSE) to demonstrate that even at very low sampling rates, audio reconstruction can be successful when enough nodes participate. In fact, when sufficient sensors participate, typically four or more, the result is an intelligible reproduction of human voice [7].

MEDIA ACCESS CONTROL AND ROUTING

In DistressNet, communications are critical. Information produced by teams and sensors enables accurate situational awareness permitting reactive precise responses to developing situations. Decisions made by C2 elements must be rapidly disseminated to affected systems and areas. However, due to its complex, highly mobile nature, DistressNet suffers from two significant problems: congestion in core regions and frequently disconnected groups on the

periphery. The solutions to these problems are adaptive spectrum-aware MAC [8] and intelligent multipath routing [9].

Multichannel MAC — In heavily populated regions of DistressNet, the network may experience congestion. Traditional single-channel MAC protocols, designed to handle low-duty-cycle traffic, are insufficient. DistressNet requires an energy-efficient solution capable of distributing users and groups intelligently across available spectrum and guaranteeing throughput with low delay. Existing multichannel MAC protocols display poor energy performance, largely due to rigid scheduling and high volumes of coordination traffic between individual nodes. Complex, fragile rendezvous protocols also impact flexibility and scalability in the network, wasting energy and increasing delay, especially when traffic volumes are lighter.

To address such issues, we propose a negotiator-based multichannel MAC protocol able to manage radio spectrum access in dense networks. The protocol, depicted in Fig. 4a, uses designated negotiators to manage schedule information and coordinate on behalf of individual nodes. Negotiators monitor a default channel continuously and coordinate agreements that allow nodes to communicate.

Simulations analyzing the use of three, four, and six channels reveal good congestion management and low packet delivery latency, even at high data rates. Figure 4b considers the impact of multiple channels on latency, and notes a direct correlation between the number of channels and the delay suffered at high data rates. This multichannel MAC also offers other benefits, including overall traffic reduction, due to a reduced need for internode coordination, and reduced storage requirements on individual nodes.

Multipath Energy-Efficient Routing with Delay Guarantees — Routing within DistressNet is difficult due to its highly mobile nature and varied topology. While routing in the core is




Figure 4. Multi-channel media access control: a) multichannel multipath routing with end-to-end delay guarantees; negotiators, N1–N4 coordinate the use of three different radio channels to avoid congestion; b) average delay experienced at various packet rates using 3, 4, and 6 channels.

complex due to high degrees of connectivity and rapid arrivals and departures of nodes, entire groups on the periphery may be disconnected for extended periods of time. Energy efficiency requirements and safety-critical messaging guarantees complicate solutions. To meet this challenge, DistressNet uses combinations of two types of routing protocols.

For connected parts of the network, DistressNet uses an on-demand routing scheme to adaptively balance energy efficiency and end-to-end delay based on message prioritization. Safety-critical messages are routed with minimal delay. Less critical messages are routed with a goal of maximizing energy efficiency.

In sparser segments of the network DistressNet relies on DTN protocols to transfer multiple copies of messages between the network core and disconnected groups using mobile VehicleNet nodes. Acting as data ferries with store and forward capabilities, these mobile nodes transfer data in both directions as their routes permit. Safety-critical messages are flooded from the network core to disconnected networks using all available resources.

LOCALIZATION

Localization is critical to the proper functioning of DistressNet and common to every element. Once localized, every node provides its location to centralized situational awareness applications. However, robust accurate localization is difficult due to the sparsely populated nature of many areas within DistressNet. Sparseness impacts localization accuracy and renders most multihop algorithms ineffective. A combination of indoor and outdoor environments increases the complexity of any solution.

Reference [10] reviews commonly used localization techniques and proposes a novel composability mechanism. DistressNet extends this idea to include external localization support and enhancements using estimation techniques. Composable localization orders available localization protocols into a library and allows the

execution of multiple localization schemes. Results obtained are compared, and the most accurate location possible is selected. Within DistressNet, localization mechanisms vary depending on the constraints affecting a node.

Mobile VehicleNet and static AreaNet nodes are localized using GPS. These and other similarly localized nodes act as beacons within DistressNet. Nodes near these beacons localize themselves using ranging protocols. However, various difficult environments within DistressNet require special techniques, including frequently disconnected nodes, localization inside buildings or where radio signals are degraded, and nodes with insufficient data to self-localize. We propose two solutions to resolve these cases: fuzzy localization and centralized localization support.

Indoors, a node may be unable to find enough beacons to localize itself. Distances to unlocalized and possibly mobile neighbors must be calculated to compensate for the lack of beacons. Normal ranging techniques infer distance based on received signal strength (RSS) but multipath propagation and channel fading, exacerbated in indoor environments, greatly affect signal strength. Fuzzy logic overcomes this by training the system to accurately map signal strength to distance, even in cases where RSS is altered by the environment. Beacons cooperate to generate a set of *rules* relating signal strength to distance in the fuzzy domain. The results are used to deduce distance using a center-average defuzzification technique triggered by hearing a message broadcast by a node desiring localization. Using distances to localized beacons, a set of nonlinear equations in the fuzzy domain is solved to obtain an estimate of the node's location. Simulations reveal that mobility and a dynamically changing topology do not significantly affect accuracy.

Another solution to some difficult scenarios is external localization support, where robust servers leverage current and historical data sources to assist an unlocalized node. When the localization support server has enough data to estimate a location, it informs the node of its




Figure 5. Cut detection using a virtual superimposed grid with a base station serving as a “source node”: a) the state of a node connected to the source node; note the early convergence before a cut at iteration 65; b) a network with a superimposed virtual grid; a cut is depicted by dotted lines; c) the state of a node disconnected from the source; note the early convergence before a cut at iteration 65.

location along with a degree of confidence. Nodes localized in this way become less reliable beacons for their neighbors.

TOPOLOGY MANAGEMENT

Rapid and frequent topology changes are endemic to DistressNet. High degrees of mobility, frequent arrivals and departures, sparse node deployments, and a typically urban setting all conspire to cause frequent network reconfiguration. Critical to the proper functioning of DistressNet is an ability to detect and compensate for topology changes. In some cases, like those involving VehicleNet, the topology change is announced. As the vehicle travels, it periodically beacons its current location. A node hearing a new beacon can assume that local topology has changed. Other changes, like the simple disappearance of a peer, may be more difficult to detect. In these cases prompt, reliable detection of this silent topology change, or cut, is important. The technique of cut detection provides notification of topology changes and supports topology management.

The state-of-the-art cut detection algorithm, Distributed Source Separation Detection (DSSD) [11], recognizes cuts by maintaining a scalar state that is updated periodically through communication with its immediate neighbors. A specially designated *source node* provides the initial state and periodic updates. During normal operations, the state converges to a steady value, as shown in the first phase of Figs. 5a and 5c. However, when a cut occurs, the states of the disconnected nodes quickly decay to zero, and still connected nodes converge to a new steady state value. Hence, a node can independently detect a cut by monitoring its own state.

The DSSD algorithm can detect cuts that separate the network into multiple components of arbitrary shapes. It displays good energy efficiency and causes little congestion due to its reliance on nearest-neighbor communications. However, the geographic scale and density variations of DistressNet affect its execution. In the densely populated core, where nodes have a high degree of connectivity, frequent exchange of state messages can cause congestion. On the periphery, frequently disconnected network segments miss state updates. To solve these problems, DistressNet uses the robust energy-efficient cut detection technique (RE²-CD) [12].

Assuming that nodes are properly localized, as indicated in Fig. 5b, RE²-CD divides the network into clusters based on group membership and location. In each cluster a leader is elected and made responsible for tracking state for the group. DSSD is executed on the resulting simplified network of leaders. If a network separation occurs and any leader is disconnected, the states of the other leaders change. Upon detecting the inconsistency in the state, a leader reports this event with its locations to the C2, where the information is consolidated and the disconnected region and groups are noted.

Critical to this cut detection technique is the accuracy of information received from neighboring leaders. Wrong information can lead to a false decision on a network separation, resulting in wasted time and energy of the operation

team. Through the use of a robust outlier rejection scheme, RE²-CD provides leaders with the capability to determine if received information is reliable enough to be used in the decision process. The level of accuracy is a tunable system parameter, and higher accuracy requires more processing power and a greater number of iterations of the algorithm.

Simulations have analyzed the performance of RE²-CD and compared it to DSSD in a virtual network of 264 uniformly distributed nodes. Figure 6a compares convergence speeds and shows that RE²-CD has an almost constant convergence speed thanks to its grid topology. Figures 6b and 6c compare the message overhead of the two algorithms for a single iteration and the total number of packets required to reach convergence, respectively. RE²-CD displays good energy efficiency at both low and high network densities.

CONCLUSIONS

We propose an architecture, a set of protocols, and required applications to support disaster response. The architecture, DistressNet, is designed expressly for the alternately congested and sparse environments typical of a disaster area where conventional infrastructure is not available. Energy-efficient by nature, the network incorporates innovative distributed sensing enhanced with collaborative features designed to maximize coverage and minimize communications. A spectrum-aware multichannel MAC protocol adds another element of energy efficiency by minimizing congestion through intelligent use of available frequencies. On-demand and delay-tolerant routing serves both connected and disconnected segments of the network. Composable localization, coupled with fuzzy estimation techniques, provides accurate locations and supports situational awareness. Industry-standard networking allows hardware- and operating-system-agnostic messaging, supports heterogeneity at all levels, and ensures extensibility. All of these elements combine to provide first responders with detailed situational awareness suitable for mission-critical decision support.

REFERENCES

- [1] A. Arora et al., "ExScal: Elements of an Extreme Scale Wireless Sensor Network," *Proc. 11th IEEE Int'l. Conf. Embedded Real-Time Comp. Sys. Apps.*, Aug. 2005, pp. 102–8.
- [2] T. He et al., "VigilNet: An Integrated Sensor Network System for Energy-Efficient Surveillance," *ACM Trans. Sensor Net.*, vol. 2, no. 1, 2006, pp. 1–38.
- [3] A. Wood et al., "Context-Aware Wireless Sensor Networks for Assisted Living and Residential Monitoring," *IEEE Network*, vol. 22, no. 4, July 2008, pp. 26–33.
- [4] Dept. Homeland Security, *National Incident Management System*, Mar. 2004.
- [5] K. Lorincz et al., "Sensor Networks for Emergency Response: Challenges and Opportunities," *IEEE Pervasive Comp.*, vol. 3, no. 4, Oct. 2005, pp. 16–23.
- [6] J. Burgess et al., "MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks," *Proc. IEEE INFOCOM*, Apr. 2006, pp. 1–11.
- [7] A. Pazarloglou, R. Stoleru, and R. Gutierrez-Osuna, "High-Resolution Speech Signal Reconstruction in Wireless Sensor Networks," *Proc. 6th IEEE CCNC*, 2009, pp. 1–5.
- [8] W. Zhou and R. Stoleru, "Towards Higher Throughput and Energy Efficiency in Dense Wireless Ad Hoc and Sensor Networks," *Proc. ACM Symp. Applied Comp.*, 2010.




Figure 6. Cut detection simulation results in networks of various densities: a) the number of iterations required to achieve convergence; b) the number of packets transmitted in a single iteration of the algorithm; c) the number of packets required to achieve convergence.

- [9] M. Won *et al.*, "Towards Energy-Efficiency with Delay Guarantees in Ad Hoc Networks," tech. rep., Texas A&M Univ., 2009.
- [10] R. Stoleru, J. Stankovic, and S. H. Son, "On Compossibility of Localization Protocols for Wireless Sensor Networks," *IEEE Network*, vol. 22, no. 4, July 2008, pp. 21–25.
- [11] P. Barooah, "Distributed Cut Detection in Sensor Networks," *Proc. IEEE Conf. Decision and Control*, 2008, pp. 1097–1102.
- [12] M. Won, S. M. George, and R. Stoleru, "RE2-CD: Robust and Energy Efficient Cut Detection in Wireless Sensor Networks." *Proc. 4th Int'l. Conf. Wireless Algorithms, Sys., and Applications*, LNCS, Springer-Verlag, Berlin, vol. 5682, pp. 80–93.

BIOGRAPHIES

STEPHEN M. GEORGE (smgeorge@cse.tamu.edu) is pursuing a Ph.D. in the Department of Computer Science and Engineering at Texas A&M University. His research interests include security and adaptive behavior in large-scale wireless sensor networks, particularly in disaster response and military applications. He received an M.S. in software engineering from Southern Methodist University.

WEI ZHOU (zwei@cse.tamu.edu) received his B.S. degree from Peking University in electrical engineering in 2007. Since then he has been a Ph.D. candidate in the Department of Computer Science and Engineering at Texas A&M University. He works on networking protocols and applications in ad hoc and sensor networks. His research interests include MAC protocols in multi-channel environments, routing protocols for mobile networks, and cognitive radio applications.

HARSHAVARDHAN CHENJI (cjh@cse.tamu.edu) obtained his M.S. degree in computer engineering from Texas A&M University in 2009 and his B.Tech. degree from the National Institute of Technology Karnataka, India, in 2007 specializing in electrical and electronics engineering. He is currently a Ph.D. candidate at Texas A&M University under Dr. Radu Stoleru. His research interests include wireless sensor networks and localization techniques for resource constrained devices.

MYOUNGGYU WON (mgwon@cse.tamu.edu) received his B.E. degree with honors from Sogang University, Seoul, Korea. He is currently pursuing his Ph.D. degree in the Department of Computer Science and Engineering at Texas A&M University. His research interests include topology control, energy-efficient routing protocols, and distributed computing in wireless ad hoc and sensor networks.

YONG OH LEE (hisfy205@tamu.edu) received B.S. and M.S. degrees from Yonsei University, Seoul, Korea, in 2005 and 2007, respectively. Since fall 2007 he has been at Texas A&M University where he is a research assistant working toward a Ph.D. degree in electrical and computer engineering. His research interest is QoS routing protocols in wired and wireless networks.

ANDRIA PAZARLOGLOU (andria@cse.tamu.edu) received a Diploma in computer and communication engineering in 2005 from the University of Thessaly, Volos, Greece. She received an M.C.S. from Texas A&M University in 2009, where she was a research assistant working on parallel systems and wireless sensor networks. Her research interests are in the field of wireless sensor networks with emphasis on high-data-rate transmission and data compression.

RADU STOLERU (stoleru@cse.tamu.edu) is an assistant professor in the Department of Computer Science and Engineering at Texas A&M University. His research interests are in deeply embedded wireless sensor systems, distributed systems, embedded computing, and computer networking. He has authored over 35 papers and won the Outstanding Graduate Student Research Award from the Department of Computer Science, University of Virginia in 2007. He received a Ph.D. in computer science from the University of Virginia in 2007.

PRABIR BAROOAH (pbarooah@ufl.edu) is an assistant professor in the Department of Mechanical and Aerospace Engineering at the University of Florida. He received his Ph.D. degree in electrical and computer engineering in 2007 from the University of California, Santa Barbara. From 1999 to 2002 he was a research engineer at United Technologies Research Center, East Hartford, Connecticut. He received his M.S. degree in mechanical engineering from the University of Delaware in 1999.