

© 2008 Michael David LeMay

DEPENDABLE EMERGENCY-RESPONSE NETWORKING BASED
ON RETASKABLE NETWORK INFRASTRUCTURES

BY

MICHAEL DAVID LEMAY

B.S., University of Wisconsin-Eau Claire, 2005

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2008

Urbana, Illinois

Adviser:

Carl A. Gunter

Abstract

Data networking can aid disaster recovery efforts by allowing victims to contact rescuers, rescuers to communicate among themselves, and concerned friends and relatives to contact victims. Unfortunately, conventional data networks themselves are often destroyed by disasters, currently rendering these services unavailable. My thesis is that heterogeneous robust subnetworks that manage to survive a disaster can be enhanced and dynamically retasked to form an Emergency-Response Network (ERN) using techniques from mobile ad-hoc networks. In this dissertation, we discuss the challenges that arise in such applications, with particular attention being paid to security challenges. We describe specific solutions to the challenges of *emergency detection*, *platform support*, and *topology planning and assessment*, relying on the philosophy espoused by the pioneers of the Internet, that protocols and related mechanisms should be as simple as possible, to make it easy to develop correct and interoperable implementations and resist the accumulation of gold-plated requirements that restrict the applicability of the mechanisms. Then, we demonstrate several emergency-response applications running on a prototype ERN based on ZigBee and UDP/IP, and explain how such ERN applications could be deployed on realistic networks.

In the service of the Lord and my Savior, Jesus Christ, and to my loving family.

Acknowledgments

We received valuable guidance from Chandra Chekuri on the use of his network optimization algorithms. We would like to thank the TCIP Center team for their feedback on this work. This work was supported in part by NSF CNS05-5170 CNS05-09268 CNS05-24695, NSF CNS 07-16421, ONR N00014-04-1-0562 N00014-02-1-0715, DHS 2006-CS-001-000001 and a grant from the MacArthur Foundation. Michael LeMay was supported by an NDSEG fellowship from the AFOSR. The views expressed are those of the authors only.

Table of Contents

List of Figures	vii
List of Abbreviations	viii
1 Introduction	1
2 Background	6
2.1 Advanced Meters, Surveillance Camera Meshes, and Building Automation Systems	6
2.2 802.15.4 Mesh Networks	10
3 Related Work	13
4 Retaskable ERN Design	19
4.1 Network Access Control and Emergency Detection	19
4.1.1 Network Access Control	19
4.1.2 Role Management	20
4.1.3 Emergency Indications	21
4.1.4 Unintended Network Interactions	28
4.2 Platform Support	28
4.2.1 Network and Device Availability	28
4.2.2 Universal Network Access	30
4.2.3 Routing	32
4.2.4 Emergency-Response Applications	34
4.3 Network Readiness Planning and Assurance	37
4.3.1 Readiness Assessment Algorithms	38
4.3.2 Example Scenario	43
5 ERN Protocol Specification	45
5.1 Network and Transport Layers	45
5.2 Application Layer	47
5.2.1 Routing	47
5.2.2 ERN Applications	50
5.2.3 QoS	53
6 Implementation and Testbed	57
6.1 Functional Prototype	57
6.1.1 Overview	57
6.1.2 Prototype ERN Applications	59
6.2 Testbed Characteristics	67
7 Evaluation	69
8 Conclusion	73

References	74
Author's Biography	78

List of Figures

1.1	Examples of common networks that typically remain unconnected, but that could be retasked and used to provide ERN services.	3
2.1	Advanced metering network interactions.	7
4.1	Upper bound on emergency indication entropy threshold that will permit non-malicious nodes to resist attacks by malicious nodes.	27
4.2	Smart thermostat enhanced with ERN functionality.	35
4.3	Initial network provisioning graph.	38
4.4	Simplified network provisioning graph with sources (light gray), sinks (white), and routers (dark gray) identified. Also includes provisional links in blue.	39
5.1	ERN network layers and their relationship to the layers of the underlying networks used to form the ERN.	45
5.2	ERN network and transport layer frames, plus the segment header.	46
5.3	Acknowledgment and routing frame formats.	48
5.4	High-level application ERN frames.	52
5.5	Application-level ERN frames for node information lookups.	54
5.6	Quality-of-Service control frames.	55
6.1	Configuration utility for setting parameters of Maxstream XBee Series 2 ZigBee modems.	58
6.2	Network audit console.	59
6.3	Main ERN application panel.	61
6.4	Node information database management dialog.	61
6.5	Node information database registration dialog.	62
6.6	Node information database query dialog.	62
6.7	Dialog to search for nodes by role.	62
6.8	Panic button request dialog.	63
6.9	Emergency indication dialog.	64
6.10	CAP alert reception dialog.	64
6.11	Text messaging dialog.	65
6.12	Audio messaging dialog.	65
6.13	Image transmission dialog.	66
6.14	Topology of implementation testbed.	67
7.1	Packet latency and throughput observed during image transfer experiment across one UDP/IP hop.	70
7.2	Packet latency and throughput observed during image transfer experiment across one ZigBee hop and then one UDP/IP hop.	70
7.3	Packet latency and throughput observed during image transfer experiment across one ZigBee hop.	71

List of Abbreviations

AMI	Advanced Metering Infrastructure
BAS	Building Automation System
CAP	Common Alerting Protocol
CCM	Counter with CBC-MAC
DER	Distributed Energy Resource
DR	Demand Response
DTN	Delay-Tolerant Network
EAS	Emergency Alert Service
ECDH	Elliptic-Curve Diffie Hellman
ECN	Explicit Congestion Notification
ERN	Emergency-Response Network
ESP	Energy Service Provider
HAN	Home-Area Network
HVAC	Heating, Ventilation, and Air Conditioning
MDMA	Meter Data Management Agency
QoS	Quality-of-Service
TTL	Time-To-Live
VAN	Vehicle-Area Network
VoIP	Voice-over-IP
WMN	Wireless Mesh Network

1 Introduction

Recent events have demonstrated the susceptibility of conventional network infrastructures to both man-made and natural disasters. The attacks on September 11th, 2001 disrupted many communication channels that were routed through the World Trade Center, and the mass panic that ensued also caused the telephone switching network to collapse [BOR02]. Even more significant disruptions to communication channels occurred when Hurricane Katrina rendered most of the infrastructure components within its wake partially or completely inoperable [CH06]. This caused great difficulties for both the victims of Katrina and those who were working to save them.

Network connectivity is very important in the aftermath of a disaster, as it can be used by victims and rescuers to communicate among themselves, send messages out to unaffected areas, and receive critical information from external sources. In the rescue operation that followed Katrina, it would have been helpful to rescuers if victims had been able to communicate their locations to rescuers, rather than forcing them to search every house. Analysis of the Kobe earthquake also cited a lack of communication as a cause for delayed emergency-response actions and a mis-direction of resources to areas that had less urgent needs than other areas [TG97]. Thus, it is clear that resilient data networks could have provided great benefits in the aftermath of this disaster, and the many others like it that occur.

After Katrina, the only significant, operational network in New Orleans was a Wireless Mesh Network (WMN) used to transport data from security cameras in the city [Gre06]. City officials used this network to provide the services normally provided by other networks, such as voice messaging (VoIP) and general police communications. A number of other major cities are now planning to deploy dedicated mesh networks to improve the robustness of the information infrastructure used by government personnel. Independently, many commercial mesh networks are being deployed for various purposes. For example, traditional electric meters are being replaced with advanced

meters that have computational capabilities and are often connected to the Meter Data Management Agency (MDMA) using mesh networking [Res05]. Buildings are also being enhanced with mesh networks for building automation [Ega05]. Vehicle-Area Networks (VANs) have also generated a significant amount of interest recently, and seem poised for deployment in the future [HBZ⁺06].

Dedicated emergency-response wireless networks can be reactively deployed in areas affected by emergencies [MB02], and some have even been installed in certain critical locales preemptively, so that they can immediately respond to future emergencies [Gov06]. However, like any other infrastructure improvement, it is often expensive to deploy such networks on a wide scale. Even if a mesh network is deployed, the number of nodes it contains must be based on the amount of functionality and value it provides. If a network is only available for use during emergencies, the level of value it provides may be relatively low compared to existing commercial networks, ultimately causing it to be comparatively small and thus limiting its coverage during emergencies. It can be very expensive to reactively deploy ERNs in disaster zones, since aerial vehicles and other expensive deployment mechanisms may be required. Large, dedicated Emergency-Response Networks (ERNs) are even more difficult to justify, particularly if there is a low probability of a disaster happening in some covered area [HKH05]. Ideally, all significantly-populated areas should be covered by ERNs since disasters can occur anywhere, so other solutions are required.

Additionally, as a general principle, rarely-used systems tend to be poorly maintained and less likely to function properly when required. This suggests that the best approach is to retask existing networks for emergency-response purposes in times of disaster. In such conditions, the primary purpose of the network may not be necessary anyway, as is the case with advanced electric meters that are not required to transmit measurements when a power outage has occurred. In fact, the economic interests of the network owner may be furthered by supporting emergency response if their revenues are tied to activities in the affected region, since the availability of ERN may permit the affected region to recover more quickly and return to normal business. It may be possible to modify existing networks to provide this service, but we also discuss requirements for future networks that will ensure they can be retasked to support emergency communications when necessary.

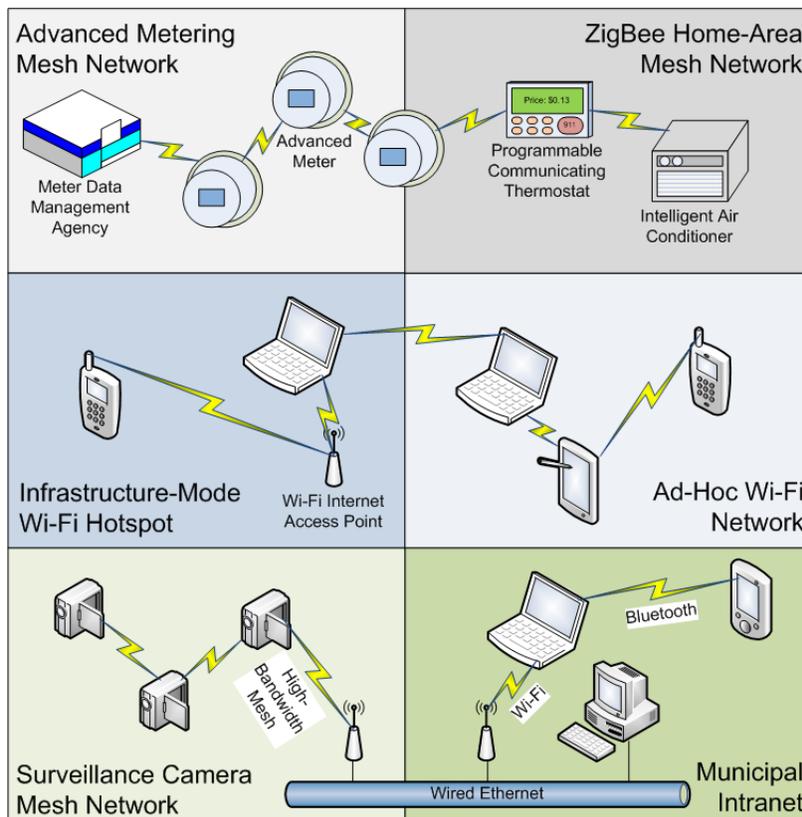


Figure 1.1: Examples of common networks that typically remain unconnected, but that could be retasked and used to provide ERN services.

In this dissertation, the overall challenge that we address is how to use various types of pre-installed networks, in particular advanced electric metering networks, Wi-Fi mesh and infrastructure networks, GSM cellphone networks, and Ethernet wired networks (depicted in Figure 1.1), to provide emergency communications while also ensuring that unauthorized parties are unable to abuse those networks in normal circumstances, and are unable to compromise any security-critical networks even during emergency conditions. There are at least three primary considerations that lead to a solution to this problem: *emergency detection*, *platform support*, and *topology planning and assessment*. First, it is necessary to establish the policies and mechanisms by which devices within the network will detect the presence of valid emergency conditions and adapt to them. Second, it may be necessary to have special emergency-response hardware and software platform support provided by devices both internal and external to the network. In particular, we show that ad-hoc network routing protocols provide a robust mechanism for supporting emergency-response networking on heterogeneous networks with very different underlying routing protocols. Third, it may be desirable to anticipate the support that will be provided by the fixed network topology itself, to ensure that ERN services are available regardless of the presence or absence of mobile nodes that may provide ad-hoc infrastructure enhancements. Of course, it can be beneficial for an ERN to permit mobile nodes to join the network and offer routing services to extend its coverage and bandwidth, but such nodes can not necessarily be relied upon as emergency service providers, since their locations may be unpredictable.

We propose solutions to all three of the aspects of ERN discussed above and implement a prototype retaskable ERN that uses ZigBee and IP networks. We analyze the performance of our prototype and discuss some of the challenges that arise when dealing with low-bandwidth mesh networks. In brief, we propose both centralized and distributed emergency detection techniques, the use of simple application-level protocols to provide emergency services, and MANET-oriented routing protocols for network operation. We propose a role-based QoS scheme to ensure that legitimate disaster victims and rescuers are given priority over other network users. Finally, we propose that an adapted buy-at-bulk network provisioning algorithm be used to assess the readiness and cost effectiveness of network infrastructure for ERN retasking. It can also be used to plan and optimize infrastructure enhancements.

This dissertation comprises eight chapters. The second chapter provides background on the major concepts referred to in the rest of the dissertation. The third chapter discusses related work in a variety of areas. Chapter Four outlines our solutions to the challenges discussed above, culminating in a complete design for a retaskable ERN. The details of our ERN protocols are contained in Chapter Five. We describe our prototype implementation of that design in Chapter Six, and experimentally evaluate it in Chapter Seven. Finally, we conclude in Chapter Eight.

2 Background

This section provides background information on the important components used in our system.

2.1 Advanced Meters, Surveillance Camera Meshes, and Building Automation Systems

We use Advanced Metering Infrastructure (AMI), surveillance camera mesh networks, wired and wireless Ethernet networks, and enhanced cellphones as running examples throughout this dissertation. AMI networks are a relatively recent development, so we provide some background on them here.

Advanced meters are electric meters that have been modified so that they transmit usage information back to an MDMA using a network. Some meters also support advanced control functions. Advanced meters afford a number of potential advantages to ESPs, their customers, and many other entities [Con04]: *1) Customer control:* Customers gain access to information on their current energy usage and real-time electricity prices. *2) Demand response:* Power utilities can more effectively send control signals to advanced metering systems to curtail customer loads, either directly or in cooperation with the customer's Building Automation System (BAS). Current Demand Response (DR) schemes are typically very coarse-grained and provide marginal power savings. *3) Improved reliability:* More agile DR and Distributed Energy Resource (DER) management can improve the reliability of the distribution grid by preventing line congestion and generation overloads. These improvements will also reduce the strain on the transmission grid. *4) Simplified sub-metering:* Multiple customers can be monitored by a single meter, reducing equipment costs and maintenance burdens. In some settings, it may even be possible for an MDMA to collect readings from multiple meters in a hierarchical fashion.

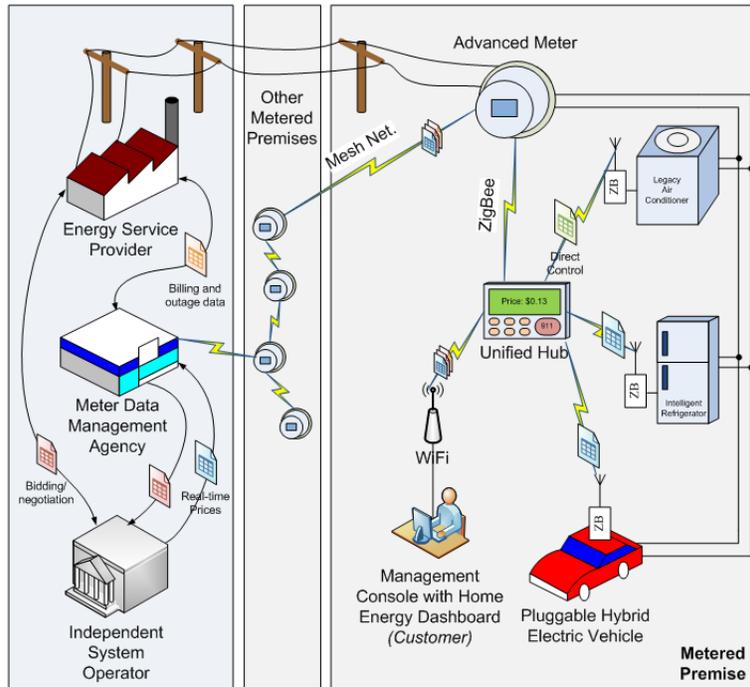


Figure 2.1: Advanced metering network interactions.

There are several distinct categories of advanced metering systems that support the functionality discussed above with varying degrees of success. The least capable systems use short-range radio networks, requiring readers to drive by in vans to read the meters. More capable systems support unidirectional fixed network communication, and the most capable systems have fully bidirectional network connections. The less capable systems are typically less expensive to deploy initially, but fully networked systems provide more economic benefits in the long run [BJR02]. Thus, we concentrate on meters with bidirectional mesh network connections throughout this dissertation. Mesh networking has started to achieve significant penetration into the AMI market, so this is a legitimate emphasis [Res05].

In Figure 2.1, we show how a bidirectional metering network could be organized. The network is divided into two main domains that are connected via a WAN link. The first domain houses the MDMA and its associated applications, such as those for analyzing metering data. The second domain comprises the metered premises, which have mesh network connections between themselves to extend the overall reach of the metering network. Each of these premises may also be equipped with a facilities LAN

containing a consumer portal, which interacts with a consumer portal application on the meter. The LAN also provides connectivity for a management console from which the customer interacts with the consumer portal, possibly using a web browser as the interface. The functionality of the consumer portal is not yet fully defined, but may permit customers to access information on their usage history, DR preferences, etc.

Individual advanced meters equipped with mesh networking interfaces do not truly fit the typical profile of either mesh routers or a mesh clients. Like routers, they occupy fixed locations and are usually connected to an ample power supply. However, they are only equipped with a single antenna for cost reasons, and execute business functions like a typical client. Normally, these characteristics would imply that the optimal method of operating an advanced metering network is to use each meter like a client. However, meters can be very widely dispersed geographically, so it may be necessary to use some meters as routers to extend network coverage away from central access points.

Additionally, modern meters are beginning to include interfaces to BASs or Home-Area Networks (HANs), so meters actually serve as network access points as well. BASs and HANs are collections of sensors and actuators that control various aspects of homes and buildings. For example, it is common for Heating, Ventilation, and Air Conditioning (HVAC) systems and lighting systems to be connected to a BAS in commercial buildings. Some buildings feature more extensive controls to manage electronic doorlocks, occupancy sensors, etc. Some modern homes are equipped with similar controls, as well, but HANs typically use different backend protocols (Z-Wave, UPB, X-10, Insteon, etc.) than are used in commercial BASs (ZigBee, BACnet, Lon-Works, etc.). The network interfaces advanced meters use to communicate with BASs and HANs are more likely to conform to a standard protocol such as 802.15.4 than the main metering networks themselves, so they may provide an ideal access point for rescuers and victims with independent ERN devices.

Increasingly, these controls are being used to automatically adjust the energy usage of buildings in response to demand or pricing indications from the Energy Service Provider (ESP) that provides electricity for the buildings. This is referred to as DR, since the power loads respond to signals from the power source to help manage the load on the electric power grid. DR programs can reduce electric costs and even prevent

blackouts in extreme cases. These benefits are driving the adoption of automation even in small residences, since the aggregated effects of many small participants reducing their demand can have a significant influence on the health of the grid. Thus, mesh networks serving these systems should soon be widespread in both commercial and residential areas. For example, Southern California Edison recently filed a plan to deploy 5.3 million residential meters [sce07]. If these networks were adapted to support ERN services, they could greatly increase coverage with relatively small additional investments.

As mentioned above, in at least one major disaster (Katrina) mesh networks were the only networks to survive the disaster. One drawback of past mesh networks was their relatively small scale. The mesh network in New Orleans was sufficient for limited (business) VoIP traffic and police communications, but was unable to provide service to ordinary citizens. In contrast, it can reasonably be expected that advanced metering mesh networks will be deployed widely within the next decade. These networks will dwarf existing meshes and although they typically use networking technologies with much less bandwidth than current mesh networks, many more nodes can potentially be supported since more routes will be available between them.

AMI and surveillance camera networks share several characteristics, but are also quite different from each other. Both surveillance camera and AMI networks are primarily local-area, and thus are unable to transfer messages long distances without additional infrastructure. Techniques for establishing emergency-response backbones are presented in [MB02], and delay-tolerant networking may also be useful in certain situations [HABR05]. Enhanced cellphones could be particularly helpful in this regard, because their radios are relatively powerful and could greatly extend the coverage of an ERN. What we refer to as enhanced cellphones are cellphones that either contain additional network interfaces to permit them to interact with other networks during emergencies, or whose standard network interfaces have been modified to permit them to communicate directly with other cellphones and network devices without interacting with cellphone towers, which are often particularly vulnerable during emergencies.

The three primary network types we have discussed also offer vastly different user interfaces. The commercial BASs connected to industrial advanced meters sometimes contain display panels accessible to building occupants, while residential HANs of-

ten contain smart thermostats. These display panels and smart thermostats could be enhanced to provide a convenient user interface by incorporating an emergency notification button that occupants press to indicate the presence of emergency conditions and request help from rescuers in their area, for example. Other devices, such as automated lighting, could be used to bring residents' attention to pre-emergency alerts. Of course, cellphones include very sophisticated integrated user interface capabilities that could be similarly enhanced. On the other hand, surveillance camera networks are often connected to corporate intranets, to provide access to imagery, and thus are only equipped with a centralized user interface.

One final distinguishing characteristic of cellphones is their mobility. Advanced meters and surveillance cameras have fixed locations, so their connectivity topologies may be more predictable and reliable than those of mobile cellphones, although fixed wireless networks may still be affected by transient atmospheric conditions and terrain changes that may occur during certain disasters.

2.2 802.15.4 Mesh Networks

In this section we review the important properties of a mesh networking technology that has achieved prominence in advanced metering and building automation applications and that we use in our prototype: ZigBee [Zig06].

ZigBee is a wireless protocol stack for low-rate wireless personal-area networks [Zig06]. It is distinctive because of its integral support for mesh networking and a strong emphasis on protocol simplicity to enable inexpensive, highly power-efficient implementations. It is built upon the IEEE 802.15.4 MAC layer [LAN03], which commonly operates on the 900MHz or 2.4GHz frequency bands. The theoretical bandwidth limitation of radios operating at 2.4GHz is 250kbps, but we have never achieved usable single-duplex data rates significantly exceeding 60kbps in our experiments. 802.15.4 relies on 16-bit network addresses for ordinary transmissions, so there is a theoretical limit of over 65 thousand nodes per network.

This standard addresses a specific set of requirements exhibited by many sensors and control devices [Kin03]. Most importantly, these devices must be extremely power efficient. In fact, 802.15.4 is designed to permit a device to run for up to two years

between battery charges or replacements. It accomplishes this primarily by reducing the range (50-100m outdoors) and bandwidth (up to 250kbps) of each node in the network compared to other network standards such as Wi-Fi, although any sufficiently dense network should still permit multi-hop communications between distant nodes. 802.15.4 also defines a number of aggressive sleep modes that permit nodes to only wake up occasionally to transmit short bursts of data, and then quickly return to a power-conserving sleep state.

Another significant requirement of many sensor networks is that messages experience low latency, since many messages in such applications have strict realtime requirements. 802.15.4 can be operated in a special mode to support realtime applications. This mode is called “beaconing mode” and requires a network coordinator to transmit beacon messages periodically that demarcate “superframes,” each of which comprises 16 identical time slots. Some number of these slots can be reserved for particular nodes so that those nodes are not required to contend for bandwidth. All other nodes with data to send contend for the remaining slots. Unfortunately, the devices used in our prototype do not support beaconing mode, so we were unable to evaluate it.

ZigBee provides basic security operations to prevent nodes outside a particular network from eavesdropping on communications within that network or injecting their own traffic [Zig05]. It uses the 128-bit AES encryption primitive to provide confidentiality and authentication for data. It can use a single key to encrypt all communications within the network, or individual pairs of nodes can share a key to protect their communications from other nodes in the network. Keys are managed by a single “trust center” in the system, but the policies for transferring keys to new devices that join the network are left unspecified.

If keys are transferred over the air it is possible for a malicious node to capture the key as it is being transmitted. Thus, the ZigBee Alliance recommends transferring keys out-of-band before connecting devices to each other. Due to the eavesdropping threats, it is best to avoid transferring the main network key to new nodes when entering emergency mode, so as not to reveal the network key to the arbitrary nodes that are allowed to join the network in that mode. Instead, a new key should be used to communicate with emergency nodes, while retaining the original key for inter-meter

communications. Going beyond basic encryption, freshness counters prevent replay attacks. Optionally, integrity checks of varying strength may also be included in each message to prevent undetected modifications. Of course, it is possible to disable all of these security features if they are not needed. Unfortunately, the radios used in our experiments lack support for per-link encryption, so we were unable to evaluate it.

One of the chief drawbacks of 802.15.4 is its limited range, although radios with ideal outdoor ranges of up to 1.6 km are available for use on limited channels in the 2.4GHz band, and radios with ranges of 64 km are available for the 900MHz band. Thus, 802.15.4-based metering networks are usually most suitable for dense applications such as those in East Asia, which is where ZigBee metering networks are most popular [Har05,SD05]. In less dense applications such as those in North America, proprietary mesh technologies dominate [Els05]. Of course, it will be difficult for mobile devices such as PDAs and cellphones to interact with proprietary networks. However, even these proprietary meters have begun to integrate ZigBee interfaces for performing building automation tasks [Itr06]. These interfaces can also serve as gateways between mobile devices and meters, generating a hybrid solution. Thus, we are justified in our focus on 802.15.4 and ZigBee, since it provides us with a general solution that is applicable even in these adverse circumstances.

3 Related Work

In the United States, the existing Emergency Alert Service (EAS) is used to deliver emergency notifications to citizens over radio and television stations, and using audible sirens mounted in centralized locations [Bot03, Moo06]. Recently, the limitations of this system have become apparent. Most importantly, the system is relatively underutilized. For example, EAS was never invoked on September 11, 2001, even though officials believed that New York and the rest of the country may have been experiencing the first stages of a large-scale assault. This seems to indicate that officials did not believe the system would have effectively alerted citizens, as it was originally intended to do. Perhaps they believed that standard television and radio coverage of the events was more effective than an austere EAS warning.

Whether or not this feeling is generally justified, the system does have severe weaknesses. The EAS is arranged hierarchically, so that a few large radio stations broadcast alerts to smaller stations that continue propagating the messages down the chain of stations with decreasing size. Thus, if some of the top-level stations were destroyed, the EAS alerts would never be transmitted to their subordinate stations. Another unrelated weakness is that EAS alerts may not actually reach citizens, even if they are transmitted. Those that are deaf and hard-of-hearing are unlikely to hear sirens, and in the past families have been killed in their homes because they slept through alerts transmitted late at night or early in the morning [Ebe04].

The FCC and DHS are currently enhancing the EAS to use more adaptable and pervasive notification devices such as cellphones, digital television sets that can be remotely activated, and of course Internet-connected computers [Was08, Ass06]. The critics of this proposal have made the observation that these devices are often dependent on non-resilient infrastructures that may be compromised before the alerts are delivered. In contrast, traditional radio stations may be able to run for several days on batteries or generators [Ebe04].

It is possible to combine the best aspects of both the new and old EAS systems using the devices in ERNs, particularly advanced meters and BASs. For example, the BAS of a deaf individual could flash a set of lights in the residence to attract the attention of the occupant to a message display. The actual alert message could be displayed on the occupant's computer or mobile device, or even on the small LCD display integrated into most advanced meters. If the BAS is unable to operate because of a power outage the occupant could be instructed to monitor the meter display directly for alerts, if accessible. In the following chapters, we describe the system components that are necessary to support such applications.

The EAS provides only rudimentary alert services, which are a small part of a comprehensive disaster-preparedness plan. Currently, the predominant advanced disaster-preparedness practice in major cities is to deploy dedicated mesh networks based on existing networking techniques that are capable of being used after a disaster has occurred. For example, Beijing's Public Security Bureau is deploying a large mesh network in preparation for the 2008 Olympics [Gov06]. This network is intended to carry emergency communications (including voice communications) and surveillance data. Other cities are embarking on similar projects, as can be seen on Strix Systems' press release listing [Str]. These networks are usually constructed primarily for government communications, although some of them have been harnessed for other purposes as well. Of course, only a small minority of cities are equipped with advanced ERNs, and rest of them rely on proprietary radios to support communications between emergency responders, and are inaccessible to disaster victims and other interested parties. Our proposed approach to providing ERN services rectifies this limitation.

Many basic and advanced ERNs are not interoperable with each other. To address this problem, the US government has recently redoubled its efforts to develop an interoperable ERN that can be used by all public safety personnel on the 700MHz band that will soon be abandoned after the conversion to analog broadcast television [Com], along with other bands that have already been allocated to public safety communications. A brief overview of these efforts and others like it is available in [Fau07]. It highlights the possible tragic consequences of incompatibilities between first-responder communications systems, and then outlines the actual interoperability requirements that first responders have. For example, it would have improved the safety at the site of

the Columbine massacre if the responding police officers had possessed interoperable radios, since those used by the thirty-nine responding agencies were all incompatible with each other. Fortunately, no additional people were killed due to confusion among the heavily-armed and stressed responders, but the rescue efforts of all the agencies involved were undoubtedly hampered by their inability to communicate efficiently. Many interoperability efforts focus on analog voice communications since that is the primary communication mechanism currently used by emergency responders, and [Fau07] argues that this is the most promising approach to pursue. In fact, speaking of recent efforts to use data networks for public safety communications, the authors go so far as to state that “Pushing one’s agenda of untried systems into the emergency communications area is foolhardy in the extreme.”

A similar analysis of the current state of public safety communications was performed in [FDW06], but quite different suggestions for improving the system were set forth. The authors propose a “network-of-networks” approach centered on an IP backbone that connects the individual networks of essentially all agencies and other networks relevant to public safety. In fact, [FDW06] recommends the use of shared networks, in which public safety personnel are given high-priority access to the network during times of emergency, but other commercial and selected governmental agencies would be granted access at other times. This suggestion is an inversion of our proposal, which is to only allow public safety personnel and disaster victims to access existing networks during times of emergency. Either approach may be more suitable in a specific installation, depending on the availability of government funding and the prior existence of suitable commercial networks. The authors cite similar reasons for using multi-purpose, or retaskable, networks as ERNs as we do, primarily the ability to share the cost of the network among multiple parties. Another benefit of either approach is the ability to use mass-produced devices, or at least only slightly-modified versions of those devices. Communication devices for non-commercial spectra are necessarily produced in lower quantities, driving up prices.

Much of [FDW06] is dedicated to a discussion of the political reasons for the dismal state of public safety communications interoperability. The government has already allocated separate segments of radio spectrum to different public safety agencies, providing little incentive for those agencies to cooperate among themselves. The authors

recommend several specific strategies that can be used to encourage interoperability, including: 1) only award additional spectrum to agencies that agree to participating in an IP-based, interoperable network; 2) offer public service spectrum to commercial companies for free if they agree to provide access to a public service broadband network when it is required; 3) incentivize public safety agencies to use commercial networks; 4) create a secondary market for public safety spectrum, so that public safety agencies can lease their allocations to companies on an interruptible basis, to help fund sophisticated, interoperable equipment. These recommendations are focused primarily on the requirements of public safety personnel communicating amongst themselves, and do not address the requirements of public safety personnel attempting to communicate with victims, or victims communicating amongst themselves. To support these additional requirements, it will be necessary to provide broader network coverage than may be feasible to construct with government and limited commercial funding, and it will also be necessary to provide ERN services that are specially designed for use between these additional classes of participants. As we show in this dissertation, re-taskable broadband or narrowband networks with specialized ERN services can help satisfy these additional requirements.

An overview of recent policies related to public safety communications is provided by [Moo05]. It points out that the current spectrum allocation in the 700MHz band is unlikely to be sufficient to support broadband networks, which suggests that re-taskable commercial networks may in fact provide greater bandwidth during emergencies than some dedicated ERNs. The authors also note that new technologies may permit a broader definition of emergency responder, so that it includes utility workers and even bystanders. We propose one such new technology and show how it effectively integrates these non-traditional “emergency agents,” as they are called by a Focus Group for the National Reliability and Interoperability Council (NRIC VII), which has helped to suggest a list of possible types of agents that will be able to participate in future emergency-response efforts.

A common approach to deploying ERNs is to develop portable mesh nodes that can be quickly deployed at the scene of a disaster. In fact, portable 802.11 mesh nodes were deployed in the New Orleans area after Katrina and reportedly routed 10,000 VoIP calls in three days [mat06]. However, there are some drawbacks to this approach that gen-

erally make it less desirable than using retaskable commercial networks. First, these portable nodes are quite expensive, historically on the order of \$500 per node. Commercial infrastructure components will be purchased by other agencies and can then be used with little additional investment (mostly limited to ERN software development and deployment). Secondly, portable nodes can only be reliably placed at the perimeter of a disaster zone, since it may be hazardous to enter the zone itself. This leaves rescuers and victims trapped within the zone with no means of communicating with the network. Expensive placement mechanisms such as airdrops can be used in such situations, but may not be cost-effective compared to our scheme. Advanced meters and other commercial networks will be widely deployed, and any undamaged networks can be used for communications deep within the disaster zone, avoiding these issues.

One significant advantage provided by fixed-location nodes over ad-hoc nodes deployed after disasters is that each node can be aware of its location without requiring expensive GPS or other hardware additions. Location-aware mesh routing protocols can optimize the routing of messages when this information is available and accurate. For example, the MESH routing protocol makes use of such information to reduce the number of blind broadcasts necessary during routing, and to efficiently recover routing information [CL00]. GPSR is another example of such a routing protocol that greedily attempts to maximize the forward progress of packets towards their destinations until it encounters a local minima, at which point it constructs a planar routing graph and routes around the faces of the graph until it can resume greedy routing [KK00]. We do not explore geo-routing further in this dissertation, but it could be a promising area of future work.

One of the most significant aspects of the proposed system is its node admission system, whereby external users are only permitted to access the network after the required router nodes have acquired sufficient confidence that an emergency has occurred. Other schemes exist for performing admission control to ad-hoc networks, and in fact they may be used in the future to improve upon our simple binary system of all-or-nothing access control. For example, SWAN uses the Explicit Congestion Notification (ECN) features of standard, best-effort MAC protocols to differentiate between different classes of traffic admitted to a network [ACVS02]. However, these schemes fail to explain how to determine which nodes should be granted access to the network.

We present a scheme to determine when to allow arbitrary nodes to access the network, and when to restrict access to just the nodes originally authorized by the network operator.

Another significant aspect of the proposed system is its application of ad-hoc networking protocols to heterogeneous networks, including wired networks. A U.S. patent (#20040141511) has been issued for a technique to bridge Bluetooth and Ethernet networks, potentially using AODV routing on the Bluetooth network. In contrast, our system uses AODV across all types of networks, since the original routing protocols of some of the networks may not operate properly in an emergency scenario, in which part of the network may have been destroyed.

4 Retaskable ERN Design

4.1 Network Access Control and Emergency Detection

4.1.1 Network Access Control

During normal operating conditions, business networks do not ordinarily route messages for arbitrary external entities. The companies that operate the networks wish to make use of their investment without granting access to parties that have not contributed to that investment. To enforce network access control, ZigBee uses the cryptographic techniques discussed in Chapter 2. In most low-power networks, a single symmetric key is used to encrypt all communications so that any node that does not possess the key is excluded from the network. Wi-Fi has similar access control schemes at its disposal. Wired Ethernets are commonly configured to perform MAC filtering, and more secure access control mechanisms may be possible.

When an emergency occurs, infrastructure providers may permit victims and rescuers to use whatever remains of their infrastructure to recover from the disaster. This can be beneficial to their revenue because it helps to speed the recovery of the area in which they do business. Thus, once emergency conditions have been recognized by a device, it must allow rescuers, victims, and other affected personnel to access the network, and perhaps activate emergency-response applications on the device. There are several obvious ways to accomplish this. First, the network could be made inaccessible for communications outside the primary functions of the network during normal conditions. Even if it is possible for a node to join the network, access to ERN services can be restricted by refusing to route most ERN messages. This can be more useful than completely restricting access to the network, because the ERN can then accept emergency indications from a larger number of nodes without allowing access to other services such as panic button broadcasts, etc. Second, access could be provided even during normal conditions, but QoS controls could be used to ensure that

communications related to the primary functions of the network are given higher priority than other communications, to prevent them from being adversely affected. During emergency conditions, it may be necessary to adjust the QoS controls, or perhaps the primary functions would simply cease and automatically be allocated less bandwidth.

Neither of these general strategies is necessarily superior in all instances, so the requirements of each network must be taken into consideration. For example, advanced meters are powered by very computationally-limited microcontrollers that are unable to coordinate complex QoS strategies. Thus, the former suggestion is probably most appropriate for such networks. On the other hand, the Wi-Fi networks that often support surveillance cameras are likely to have more powerful processors and may enable the network operator to achieve additional utility from their networks by permitting other data to be routed on the network during normal operating conditions.

Ideally, not every node should be provided with access to all services available through the network, such as service from a limited number of emergency responders. Instead, priority should be granted to traffic originating with rescuers and victims. It can be quite difficult to distinguish between different types of nodes in most emergency scenarios, but we discuss some mechanisms that could permit such distinctions in a secure manner based on the concept of centrally-certified roles.

4.1.2 Role Management

As stated in the previous section, various roles should be supported for ERN users. Users authorized as representatives of fire, police, medical, and other emergency-response organizations should be granted preferred access to the network. For this purpose, we require the construction and maintenance of a regional, national, or worldwide emergency-response authorization hierarchy, based on PKI. All emergency responders must be issued certificates encoding their attributes, so that those attributes can be used to authorize various levels of network access [PS00]. Monolithic certificates must be used because many devices in the ERN have very limited resources, and the overhead of verifying multiple signatures to accomplish a single operation could cause excessive delay and energy drain. For this reason it is also impossible to distribute large CRLs to all devices, so the certificates should have very short valid periods, so that compromised or out-of-date certificates expire and are removed from the CRL more quickly.

This requires more communication between the CA and authorized entities, but we assume that those devices are much more capable than the ERN devices in the field, so this is a reasonable tradeoff.

All ERNs must use a universal attribute naming scheme imposed by the attribute certification authority. For improved efficiency, it may be beneficial to encode attributes as fixed-length integers, since they are intended for direct use by machines.

Attribute certificates can be used to sign messages and thus prove the authorization of nodes to perform certain actions or make assertions. For example, authorization may be required to establish communications with and receive information from victims, propagate trustworthy information on the occurrence or spread of emergencies to other entities on the network, or obtain QoS privileges from the network coordinator. Each device in the ERN is free to maintain its own authorization policies. Obviously, there are benefits to maintaining similar or identical policies across the entire network, to allow emergency-response personnel to operate most effectively, so reference policies should be developed and made available to any devices that wish to use them as the basis for their own.

4.1.3 Emergency Indications

Causing nodes to accurately recognize emergency conditions can be challenging. We wish to avoid false positives, in which malicious nodes or users falsely indicate that an emergency has occurred so that they may illegitimately use ERN resources to route their own traffic. We also wish to avoid excessive latency between the onset of emergency conditions and the time at which legitimate victims or rescuers are able to use the ERN. In summary, we wish to ensure that no meters enter emergency mode when emergency conditions do not exist, while also minimizing the time between the onset of an emergency and the activation of emergency mode.

Before we propose a solution that satisfies these properties, we must state a very important assumption about the context of the system: We assume that no more than some percentage of ERN devices are malicious and would lie about the existence of emergency conditions. This percentage can be determined for individual networks, based upon how important it is to prevent unauthorized usage of the network and how useful the network is to the overall ERN. The former question must be decided on a

case-by-case basis, while the latter question can be answered using the topology planning algorithm that we discuss in Section 4.3. For example, a collection of restaurant hotspots in a particular region are very unlikely to host any sensitive applications, and may be able to provide a significant amount of connectivity in the case of an emergency. The primary effect of network abuse would be lost revenue from sign-on charges. Thus, it is acceptable and desirable for such networks to respond quickly to a few emergency indications, even if that makes it possible for a small percentage of users to collude and abuse the network. On the other hand, a government surveillance camera mesh network may be used to detect crimes in progress, so availability of camera images can be critical. Thus, the network should require relatively more emergency indications before providing ERN services, to prevent criminals from performing denial-of-service attacks on the camera imagery. However, this threshold must not be set too high if such a camera mesh is well-positioned to provide high connectivity during emergencies. Risk management methodologies may be helpful in making these decisions.

The general scheme that we propose to detect emergencies is to have each node listen for emergency indications from other nodes, and use those indications to probabilistically determine whether an emergency is in effect. Different nodes will have different probability thresholds that must be crossed before ERN services are provided, based on the considerations that were just presented. The recognized categories of emergencies are currently flood, fire, high wind, earthquake, explosion, biological, and radiation. These categories can be used to indicate most natural and man-made disasters, as well as acts of war or terrorism.

Emergency indication messages are broadcast messages that specify how confident the originating node is that each type of emergency condition exists. These messages may reach some nodes that are not within the region affected by the emergency, but those nodes may be well-positioned to provide ERN services to the nodes that are affected. Conversely, the broadcast radius is finite, and emergency indication messages may not reach all nodes that are affected by the emergency. This is a more difficult problem to address, but can be overcome by relying on secondary indications. For example, even if a tornado is not detected on some segment of the network by high-wind sensors, affected humans on that segment pressing panic buttons can still cause ERN services to come into effect. We will discuss the operation of panic buttons in

more detail later.

Only one emergency indication message from each node is accepted by other nodes. If a node issues a new indication message, it completely supersedes the previous one. All emergency indications also expire after some time. An expiration time can be suggested by the originator of each indication, but the recipient can place restrictions on that value to ensure that no entity in the network is able to artificially increase the value over a long period. Emergency response services are only provided by a particular node while that node is convinced that an emergency has occurred, using a process described below.

As mentioned previously, each node assigns a confidence factor to each emergency indication that it generates. For example, statistics for a particular region in which an advanced meter is deployed may indicate that when a power outage occurs, there is a 10% probability that a flood caused the outage, 2% probability that it was caused by a wildfire, and 50% probability that it was caused by high winds. The probabilities may be different depending upon the duration of the outage as well, so that after the outage has lasted for more than 2 hours, there is a 20% probability that it was caused by a flood, 8% probability that it was caused by a wildfire, and 30% probability that it was caused by high winds. Based on these probabilities, whenever a power outage occurs, the meter would update its internal indications with those probabilities, and also broadcast an emergency indication message to all other nodes in the area so that they could do the same. Many devices are equipped with multiple sensors that can be used to detect emergencies, and Bayesian inference or some other inference technique should be used to synthesize an appropriate indication from all those sensors [Pea88].

The panic button that we have recommended installing on PCTs can serve as a crude emergency indicator even in the absence of other sensors. It is impossible to tell what type of emergency may have caused a victim to press a panic button in the absence of any correlated external information. Thus, the PCT or other host device should prompt the user to specify what category of emergency best describes their observations. Then, an indication that the specified emergency has occurred with a probability of 100% should be generated and distributed, since it is unlikely that a victim will be able to provide a more accurate probability value when in distress, and because humans are equipped with many senses that permit them to assess and classify emer-

gency situations with great certainty. Beyond these widely-deployed types of sensors, some special nodes may even be distributed throughout high risk areas and equipped with more sophisticated sensors such as Geiger counters, chemical detectors, or blast sensors.

Upon receiving emergency indications, each node (agent) a_i must re-evaluate its confidence that an emergency has occurred, and change the activation state of its ERN services accordingly. The set of all agents is denoted as A . The evaluation process is repeated for each type of emergency (fire, flood, etc.). Following the inspiration of [DKG⁺04], each indication o for disaster category d received by agent a_i from agent a_j is weighted according to the following schema:

$$w_{ijo} = T_{ij} \cdot \left(\frac{3}{4} + \frac{(t_{expire} - t)}{4(t_{expire} - t_{issue})} \right),$$

where T_{ij} is the trust that agent a_i vests in agent a_j , t is the current time, and t_{issue} and t_{expire} are the times at which the indication was issued and set to expire, respectively. The weighting of each indication decreases linearly from the time of issue until the time of expiry, at which time its weight is 75% of its original weight. Each indication o is a tuple: $o = \langle a, d, c, t_{issue}, t_{expire} \rangle$, where a is the originator of the indication, d is the disaster category, and c is the confidence level that the object's originator has that a disaster of type d has occurred. The confidence c can range from 0 to 1, with values less than 0.5 indicating that the node is confident that the specified disaster conditions do not in fact exist (negative indications). Such indications are only sent from a node when that node receives a positive indication for the same category of disaster from another node and is confident that such a disaster is not underway, or has no information indicating that such a disaster is underway, in which case it would issue an indication with a confidence level of 0. It is important that nodes issue negative and neutral indications to ensure that a few malicious nodes are unable to easily manipulate the system into activating ERN services.

Authenticated emergency indications or cancellations from the network operator's central station are considered to be completely trustworthy, but may not be available during emergencies. Trust values are assigned to other roles in a network-specific fashion. It is most prudent to assign higher trust values to specially-deployed emergency sensors that are physically protected, and lower trust values to advanced meters and

other objects that are more readily accessible and equipped with less authoritative sensors. Thus, trust values are in fact assigned as follows:

$$T_{ij} = \max_{r \in \text{roles}(j)} T_{ir},$$

where $\text{roles}(j) : A \rightarrow 2^R$ associates agents with their authorized roles from the set of all roles R , and T_{ir} is the trust node i places in another node based upon its possession of role r .

The set of all indications $O_{id} \subseteq O_i$ received by agent a_i for a particular category of disaster d , where O_i is the set of all indication received by agent a_i , must have at least a certain number of objects before the agent is willing to consider activating ERN services: $\gamma \leq |O_{id}|$. Thus, if ERN services have been activated, and that relation is subsequently not satisfied, ERN services will be immediately halted. If ERN services are inactive, a new indication arrives, and that relation is then satisfied for one or more disaster categories d (regardless of whether it was previously satisfied), O_{id} is evaluated for every such d until they have all been evaluated, or ERN services have been activated on account of one of them. This should actually only occur after the node has waited for an additional short time period, to prevent malicious nodes from rapidly flooding the network with positive emergency indications without giving non-malicious nodes time to respond with negative indications.

ERN services are activated on the basis of the evidence in O_{id} if its entropy is below some agent-specific threshold, and the weighted confidence level of those indications favors the presence of emergency conditions. Entropy is computed as follows:

$$\text{entropy}_{id} = \sum_{b \in B} -\frac{W_{id}(\{b\})}{W_{id}(B)} \log_2 \frac{W_{id}(\{b\})}{W_{id}(B)},$$

where $B = \{\text{present}, \text{absent}\}$ is the set of possible beliefs regarding the presence of the relevant emergency conditions, and $W_{id}(S)$ is defined as follows:

$$W_{id}(S) = \sum_{o = \langle a_j, d, c, t_{\text{issue}}, t_{\text{expire}} \rangle \in O_{id}, b \in S} w_{ijo} \cdot \text{prob}(c, b),$$

where:

$$prob(c, b) = \begin{cases} c & \text{if } b = \textit{present} \\ (1 - c) & \text{if } b = \textit{absent} \end{cases}$$

The node is willing to consider activating ERN services iff the entropy is lower than some network-specific value, indicating a consensus among the indications received by the node. The entropy threshold must be lower than 1 bit, to ensure that it is impossible to encounter equal confidences that emergency conditions are in effect and not in effect. ERN services will be activated if the following relation is satisfied:

$$\frac{1}{\sum_{o \in O_{id}} w_{ijo}} \sum_{o = \langle a_j, d, c, t_{issue}, t_{expire} \rangle \in O_{id}} w_{ijo} \cdot c > 0.5$$

Now, we briefly analyze the security properties provided by this formulation. The security objective we are attempting to preserve is that a small number of malicious nodes should be unable to cause a non-malicious, well-connected node on the network to activate ERN services when emergency conditions do not in fact exist. We assume that the non-malicious nodes on the network have properly functioning sensors that do not indicate emergency conditions.

Let us consider the case when all nodes in the network are considered to be equally trustworthy: $T = 0.3$. To maximize their chances of success, the m attackers must issue all of their positive indications at once. We do assume that a sufficient waiting period is in place to allow n non-malicious nodes to respond with negative indications, but we simplify our analysis slightly by ignoring the reduction in the weight of the indications that occurs over time. This simplification makes the analysis more favorable for the malicious nodes if the non-malicious nodes re-issue negative indications every time they receive a positive indication.

To cause ERN services to be activated, the attackers must simultaneously achieve entropy lower than α on other nodes, and bias the activation relation towards confidence that emergency conditions are in effect. These two objectives can be achieved with the highest likelihood if the attackers use the strategy of indicating 100% confidence that a disaster has occurred. Given this optimal strategy, the following relation must be satisfied for the malicious nodes to succeed in activating ERN services:

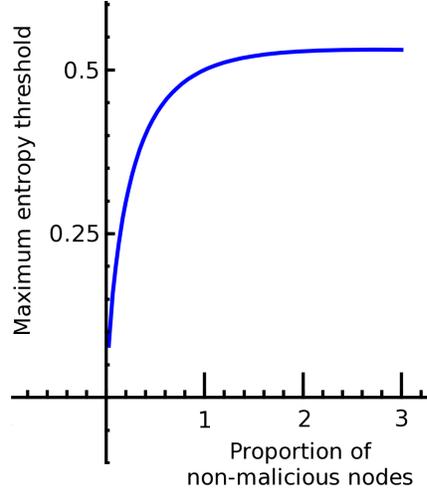


Figure 4.1: Upper bound on emergency indication entropy threshold that will permit non-malicious nodes to resist attacks by malicious nodes.

$$-\frac{n}{2(m+n)} \log_2 \left(\frac{n}{2(m+n)} \right) < \alpha$$

Figure 4.1 depicts the highest entropy thresholds that would allow the non-malicious nodes to resist the attackers. In reality, not all nodes will be equally trusted, and malicious trusted nodes will have a higher probability of overcoming a high entropy threshold. However, one factor that must be considered when computing trust values in the first place is how well protected nodes occupying trusted roles will be. Thus, it must be more difficult for attackers to compromise highly-trusted nodes. Of course, Figure 4.1 also indicates how many equally-weighted nodes must indicate full confidence that emergency conditions are in effect for ERN services to be activated. Thus, the entropy threshold must be selected very carefully, to ensure that the security requirements of the network are balanced against its availability requirements during emergencies. The properties of highly-trusted nodes that we just discussed are beneficial from both of these perspectives, since emergency indications from a few of them can overcome an entropy threshold that would otherwise require a large number of positive indications from nodes that are relatively untrusted. In our experiments later in this thesis, we use a threshold of .531, which requires more than one quarter of all equally weighted nodes to indicate full confidence that emergency conditions are in effect to activate ERN services.

4.1.4 Unintended Network Interactions

One issue that is very significant and also difficult to solve in general is that of unintended interactions between networks that provide ERN services. Ideally, every suitable network should provide ERN services, to provide maximum coverage during emergencies. However, this may introduce linkages between networks that are ordinarily air-gapped or at least separated by firewalls. For example, corporate intranets typically host resources that are inaccessible from the Internet. If such an intranet provides ERN services for both wired and wireless Ethernet users, and a nearby Internet-connected wireless Ethernet also provides ERN services, a mobile ERN node may form a linkage between those networks during emergencies. Even if such linkages are not formed, the mobile node itself may be malicious. To address such vulnerabilities, any security-critical networks that provide ERN services must ensure that firewalls are in place to ensure that only ERN traffic can be sent from unrecognized nodes, or shut-down security-critical network services during emergencies. It is also important to ensure that the ERN implementation itself is not vulnerable to attacks such as buffer overflows. This is a strong argument for simple ERN protocols that lend themselves to easily-analyzable implementations.

4.2 Platform Support

4.2.1 Network and Device Availability

Devices must be operational during emergencies to provide ERN services, and the network must not be impeded by other factors such as increased interference, etc. Of course, this implies the question of how much availability actually is required. There are at least two components to availability: *device availability* and *network availability*. Device availability is a necessary but insufficient condition for network availability. Two primary factors determine device availability: device health and device power status.

Many emergencies can adversely affect device health. For example, floods and fires can severely damage electronic circuits, nuclear explosions can produce electromagnetic pulses that destroy electronics, earthquakes can cause devices to be crushed or dislodged from protective enclosures, etc. In most of these cases, damage can be

prevented by shielding and armoring the device, as the military commonly does for critical electronics. The costs of device health preservation must be weighed against the likelihood that the benefits of emergency-response networking will be realized and the magnitude of those benefits.

Device power status during an emergency is determined by the device's power supply and the status of that power supply. Most fixed, networked devices require AC power, which is supplied by one of the infrastructures most likely to be rendered inoperative by many disasters. Thus, emergency-response devices should either support battery backup power or be connected to an AC power supply that has backup capabilities. Many modern networked devices require a minimal amount of power to maintain network connectivity, such as ZigBee devices that are designed to run for years on low-capacity batteries such as standard alkaline cells. It will be necessary to determine how much cost batteries add to the devices and weigh that cost against the benefits provided by emergency-response services. The overall reliability and availability of network edges in the presence of various hazards is a critical consideration in our network provisioning algorithm.

Current smart thermostats are often equipped with batteries that are sufficient to support their normal operations for over a year. Of course, routing an unusually large number of messages during emergency conditions will drastically shorten expected battery life. However, emergency conditions should only last for a matter of days or at most a month before survivor mortality begins to occur. Thus, the first few days of an emergency are the most critical period during which network availability must be maintained. The batteries in a smart thermostat should be sufficient to support this level of availability if the low battery indicator is aggressive enough and batteries are replaced when it lights, so that they have a sufficient charge when emergencies occur.

On the other hand, current advanced meters typically provide only enough backup power to support a few seconds of uptime, so that a last gasp outage notification can be transmitted to the MDMA. Advanced meters are a critical part of the network infrastructure because smart thermostats may not be near enough to each other to actually form a network. Thus, for advanced meters to form a useful ERN in concert with smart thermostats, they must be enhanced with additional backup power sources. A small rechargeable battery should suffice for this purpose. Standard alkaline batteries are

probably not sufficient, because meters are expected to operate in the field for up to 10 years without maintenance, which is typically outside the expected lifespan of alkaline batteries.

The necessity of ERNs varies from location-to-location, in terms of the extent of emergency-response capabilities that are required as well as the exact types of capabilities. For example, the New Orleans area is very vulnerable to flooding disasters, often trapping victims, so extensive in-home ERNs may be required in many areas. On the other hand, Oklahoma is prone to tornadoes that can destroy structures but generally do not produce floods, which are less common in most parts of Oklahoma than they are in New Orleans. People are less likely to be trapped in their homes in tornado-induced disasters, so emergency-response stations scattered throughout communities may be sufficient. Obviously, these environments affect the choice of emergency-response equipment physical protections as well. Emergency-response infrastructure in New Orleans should be protected from water damage if possible, whereas such protections would be much more difficult to justify in Oklahoma.

4.2.2 Universal Network Access

Once ERN services have been activated, the ERN must support basic types of communication normally transferred over unavailable networks. For example, the AMI network, surveillance camera network, and municipal leased lines for supporting a local government intranet serve very different purposes, and never directly interact during normal operations. However, they are often deployed in parallel, since it is typical for municipal buildings to be equipped with electric meters and be monitored by cameras. If the municipal intranet uses networking mechanisms that are vulnerable to disasters, such as wires strung between poles that can be severed by falling trees, the various local governmental agencies may become disconnected from one another during a disaster, at least from an IT perspective. Additionally, it is common for disasters to degrade or disable other communications mechanisms such as cellphone towers and landlines. Thus, the communications between governmental units may be very inefficient, perhaps relying on human messengers in cars. In spite of all this, a significant portion of the metering and camera networks may have survived if the devices are equipped with battery backup power and undamaged by the disaster. It is easy for a typical meter

or camera to be damaged by fire, falling objects, electrical surges or electromagnetic pulses, and floodwaters, but many of these hazards will be non-uniformly distributed throughout an affected region, and may only damage a portion of the devices in the region. If those devices use a self-healing mesh network, they may still be able to form fragments of network connectivity.

By routing some of their communications over these secondary networks, the local governmental units may be able to reestablish limited communications and more efficiently coordinate emergency recovery operations. Of course, networks must be provisioned to ensure that connectivity will be maintained with high probability during various disasters.

Individual devices must be properly equipped to actually access the network infrastructure during an emergency. In the case of standard protocols such as 802.11/WiFi and 802.15.4/ZigBee, the devices may reasonably be expected to include interfaces capable of directly joining the ERN. However, in the case of more obscure or non-standard protocols, different solutions may be required for different parties. Victims will not necessarily prepare for using the network in advance, but they should either have access to a device in their home that can interact with the network, or they may use the gateway solution discussed next.

It is unreasonable to expect that rescuer communication devices will include support for all network protocols in use, so gateway devices may be required to translate messages between the standard network types supported by the rescuers' devices and the ERN. Ideally, such gateway devices should be pre-installed by the network operator at strategic locations. Otherwise, rescuers may be able to install them on an as-needed basis. Of course, other creative solutions to these problems may be worthy of consideration. Victims outside buildings or in buildings not equipped with emergency communication devices must also be handled by the network. Victims possessing enhanced cellphones can directly interact with the ERN, but other cellphones are unable to do so without assistance. GPRS gateway devices running the emergency-response application on behalf of individual legacy cellphones would permit such individuals to use their cellphones to participate in the network, but other solutions may be required for certain installations. In fact, enhanced cellphones could potentially be engineered to serve as gateways for legacy cellphones.

4.2.3 Routing

Network routing is one of the most fundamental services on any network, and is particularly challenging for ERNs, due to the potential heterogeneity of the overall network. Two types of routing must be explicitly supported: Broadcast and unicast. Multicast may be useful in certain circumstances, but is generally difficult to implement. Broadcast messaging is used to propagate emergency declarations and revocations from the central authority, and emergency indications from individual nodes in the network. It is also used to propagate emergency alert messages. Unicast messaging is used to implement application communications, such as text or voice messaging between rescuers and victims.

Addressing Network addresses are fundamental to routing. Different types of underlying networks that can be interconnected by ERNs use incompatible addressing schemes, so it is necessary to use an addressing scheme that is compatible with as large a number of possible underlying networks as is feasible. We propose layering short addresses on top of the underlying networks. Disasters should not cover a very large area, so 32-bit addresses should provide a sufficient address space and not produce an inordinate amount of overhead. We refer to such an address as an “ERN address.”

Address assignment in ad-hoc networks has been widely investigated, and a particular solution to the problem is to randomly select an address, issue a routing request for that address to see if it already exists in the network, and then to continue monitoring the network to see if the address comes into use by multiple nodes at any later time [Vai02]. We adopt a modified form of this solution in our system, which we describe below.

First, a node must join whatever networks are appropriate for it and obtain native addresses on those networks using any available mechanism. Then, the node must generate a random ERN address. To determine whether the address is already in use, the node must issue a routing request for the address, using the proposed address as both the source and destination address. The request is broadcast, so it also contains a TTL value to prevent the formation of routing loops. Then, the request is broadcast to all nodes within range. Those nodes check their routing tables to determine whether they are aware of the address in question. If so, they send the request directly to the

node already using the address, so that it can generate an authenticated response with its certificate, to indicate to the new node that the address is taken. Otherwise, the node re-broadcasts the request, and records the native address of the source of the request.

If no responses arrive, the new node assumes that the address is available and acquires it. If a response does arrive, it generates a new random address and repeats the process until it is successful. A more detailed description of this process as well as the mechanism for continuously monitoring for duplicate addresses is provided in Chapter 5.

Broadcast Routing Some broadcast messages must be propagated through the entire network, whereas others should only propagate a short distance from their originators. Centralized emergency declarations and emergency alerts would fall within the first category, and distributed emergency indications in the latter. The ERN routing protocol must support both types of messages. Each message is uniquely identified by its source address and a message ID that is relatively unique to that source address. Message IDs may be reused, but only after all unused IDs have been exhausted.

Broadcast messages contain Time-To-Live (TTL) values that limit their propagation. Whenever a node receives a broadcast message, it checks its internal table to determine whether the message has already been received, and if not, rebroadcasts it and adds it to the table. Most types of networks include native broadcast functionality, which should be leveraged. In such situations, only gateway nodes should perform the actions just described.

Unicast Routing Unicast routing follows a slightly modified version of the AODV protocol [PR99]. The primary drawback of AODV compared to DSR and some other routing protocols is that it requires bi-directional links. However, this disadvantage is offset by the fact that AODV has lower overhead, since it does not require the entire route to be piggy-backed on the routing request packet. As was the case with the previous protocols, we adapt AODV to take advantage of the native routing protocols on each underlying network. We use the same routing protocol that was described in the section on address assignment for actual routing requests. One notable advantage of this approach is that AODV's drawback of requiring bidirectional links can be partially overcome on some networks. In this scenario, bidirectional links are only required be-

tween each node forwarding ERN-level routing requests and responses. Unidirectional links may be tolerated within the networks supporting those connections.

4.2.4 Emergency-Response Applications

Two distinct approaches to ERN service provision seem particularly promising. First, a suite of simple application-level protocols could be developed to accomplish whatever tasks are necessary in the aftermath of a disaster. Some Internet protocols such as SMTP have survived decades basically unchanged because of their simplicity, giving some indication that this approach to interoperability within ERNs has a high probability of success if simplicity is a driving motivation in the protocol development process. Several likely ERN services that could motivate the development of new protocols are presented in [HKH05]. We adopt this approach of simple core protocols, and describe a few that we have constructed below.

On the other hand, many of the devices in the ERN may be upgradeable, so rather than emphasizing simplicity in set-in-stone protocols that may or may not provide adequate support for changing disaster scenarios, it may be possible to simply provide an easily adaptable platform that can be reconfigured to suit the particular requirements that arise in the aftermath of a disaster. There has been considerable research on making systems extensible. For instance, *active networks* [HMA⁺99] aim to add software to network elements and there have been efforts to provide dynamically extensible platforms for sensors [FRL05]. We do not believe that this approach is currently promising for embedded systems that lack the resources to support heavyweight reconfigurable layers. In the remainder of this section we overview the ERN services we implemented in our testbed.

Of course, many other applications can be envisioned that would be useful in emergency-response scenarios. For example, an applications for tracking the vital signs of rescuers responding to a disaster site is presented in [LMFJ⁺04]. For the sake of simplicity, we eschew including support for such services in our core protocol specification, although it should be possible to define an extension mechanism in the future that could accommodate additional services.

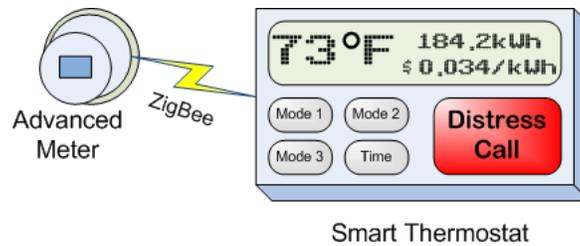


Figure 4.2: Smart thermostat enhanced with ERN functionality.

Node Lookup In chaotic emergency-response environments, nodes are unlikely to request service from specific ERN addresses. Instead, communications will often be established with nodes representing specific human individuals, or with any nodes representing a specific role. For example, a victim in a burning home is only interested in contacting some firefighter, although he is unlikely to know the ERN address of a specific firefighter. In contrast, the relative of a victim is likely to only be interested in contacting that particular victim, although this scenario is similar in that they most likely do not know the current ERN address of the victim.

In the second example, it is unlikely that the requested individual will be near the gateway to the ERN that the concerned family member's request enters. Thus, several broadcasts would likely be necessary to locate the individual using a distributed lookup mechanism. In the first example, however, it is much more likely that some entity possessing the desired role is within a few hops of the requester. Therefore, different approaches are required for these two types of lookup requests.

For specific identity lookups, one or more centralized repositories are used. The ERN address of each repository is determined dynamically, and periodically broadcast throughout the network using special emergency alerts. When a node receives such an alert, it must send its high-level identity (i.e. full legal name) and current ERN address to the repository.

For role-based lookups, a simple broadcast mechanism is used. A service request is issued, similar in format to a route request, but containing a role specification instead of an address. The service response includes the ERN address of a node serving in the requested role.

Victim Distress Calls We envision PCTs being equipped with prominent panic buttons that could be pressed by disaster victims to request contact with or in-person assistance from available rescuers, as shown in Figure 4.2. The button would ordinarily be an unlit but labeled momentary pushbutton on the face of the thermostat. When the network enters emergency response mode, the button would turn yellow. When a human presses the button, it would turn red, and the thermostat would issue a panic button request to a nearby rescuer located using the role lookup service. The request would solicit assistance from the rescuer in the form of audio or text message contact, if the thermostat is so equipped, and/or in-person assistance. The location of the thermostat should be pre-programmed or dynamically determined and included in the request if possible. Any rescuers that receive the request must respond if they are capable of providing the required assistance. Whenever the PCT receives a new rescuer response, it must send that rescuer all previously-received responses, as well as any future responses that are received. All rescuers can then examine the support pledged by other rescuers, negotiate amongst themselves, consider the requirements of other victims, and then cooperatively determine which initial offers of assistance should be honored.

Emergency Alert Broadcasts Emergency alerts are widely used across the US and other countries to notify citizens of various hazardous conditions. Such alerts can also be conveyed using ERNs. A standard format for encoding emergency alerts has been developed, and is known as the Common Alerting Protocol (CAP) [Bot03]. Other alert formats could be used, but CAP is the most widely recognized. To propagate the alert, the broadcast routing mechanism described above can be used.

Data Messaging Arbitrary data messages can be transmitted over the ERN, such as the emergency alerts just described. In our prototype, we also experiment with text, and image messaging. Different types of data are distinguished using an 8-bit code at the beginning of the message. Messages can be transmitted using unicast or broadcast messaging.

Realtime Voice Communications In addition to static messages, realtime voice communications can be used. They are actually transmitted over the network as undis-

tinguished data messages, since it is not realistic for more than one audio session to originate or arrive at a single node. This eliminates the need for session management protocols.

4.3 Network Readiness Planning and Assurance

ERNs are an important part of any comprehensive emergency-response plan, and thus should be precisely analyzed to ensure that they will be available when necessary. ERN design is fundamentally different from the design process for ordinary networks, because death or injury may occur if an ERN has inadequate provisions at the time it is required, whereas the resources allocated to ordinary networks can simply be increased after the shortage occurs with little chance of such serious consequences occurring in most cases. In this section, we outline a process for designing and assessing retaskable ERNs using a modified network provisioning algorithm originally intended for use on fixed fiber-optic networks.

First, we represent ERNs as graphs comprising data sources, data sinks, and communication links, like many networks. However, we do not assemble these into a graph in the straightforward manner. Rather, we represent network sources and sinks as vertices (where individual vertices may be both sources and sinks, or neither if they simply route messages), and represent each *possible* communication link between the vertices as an undirected edge. It may be impossible to generate a comprehensive set of vertices, because many of them will be mobile and unpredictable. Thus, in place of mobile vertices, we place virtual vertices that represent a set of physical vertices that move within some region surrounding the virtual vertex. The attributes of the virtual vertex are formulated to represent the worst-case scenario they can accommodate. Physical nodes may enter, exit, and transition between virtual regions in an unpredictable fashion. A sample graph is shown in Figure 4.3.

We can assign a number of attributes to each link, and these attributes are used to derive costs for the associated edges. Attributes of interest include the following: *cost* of constructing the link, which may be very small or zero for existing or retaskable links; the *reliability* of the link, measured as the probability that link will be operational when needed and the probability that specific hazards will occur; the *accessibility* of

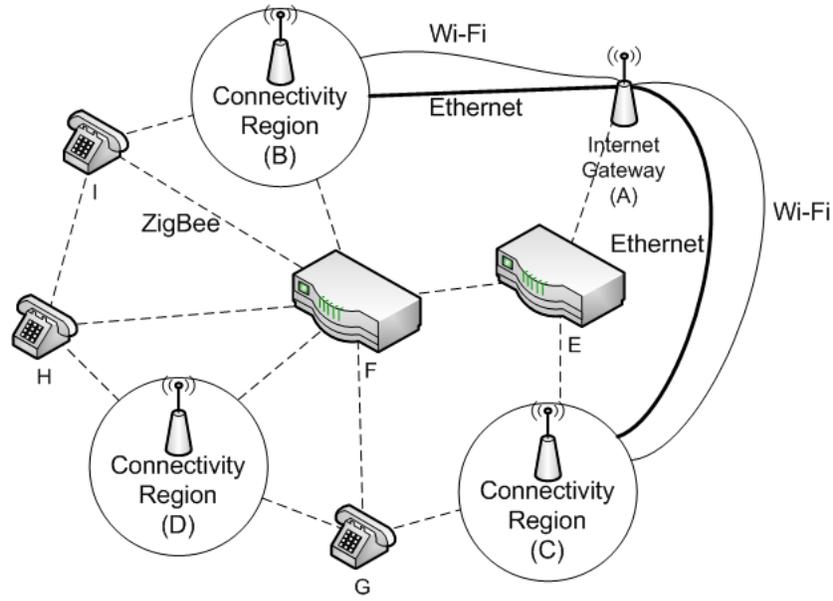


Figure 4.3: Initial network provisioning graph.

the link, measured as the likelihood that access will be given when truly needed; the *bandwidth* that can be transmitted in a unit of time over link; the *latency* for data to travel from the origin of the link to its destination.

Using these attributes, it is possible to formulate a graph model for the network. The major entities in the graph are pairs of sources and sinks, and edges that may be used in the solution, as shown in Figure 4.4. We formulate the problem as an undirected graph $G = (V, E)$ on n vertices that represents a supergraph of the network topology of any possible solution. We must also formulate a set of demand pairs $\mathcal{T} = (s_1, t_1), \dots, (s_h, t_h)$, where $s_i, t_i \in V$ and each pair is associated with a non-negative bandwidth demand d_i .

4.3.1 Readiness Assessment Algorithms

Let us now turn to the issue of how to assess the ERN readiness of a network configuration. The goal is to determine an approximately optimum network routing scheme by finding a feasible flow for all the pairs in which d_i bandwidth flow is sent from s_i to t_i while minimizing costs. Of course, monetary cost is not the only consideration in our problem, so we must map the other considerations to cost for the algorithm to be useful. The network topology we analyze is considered to be a multi-commodity topol-

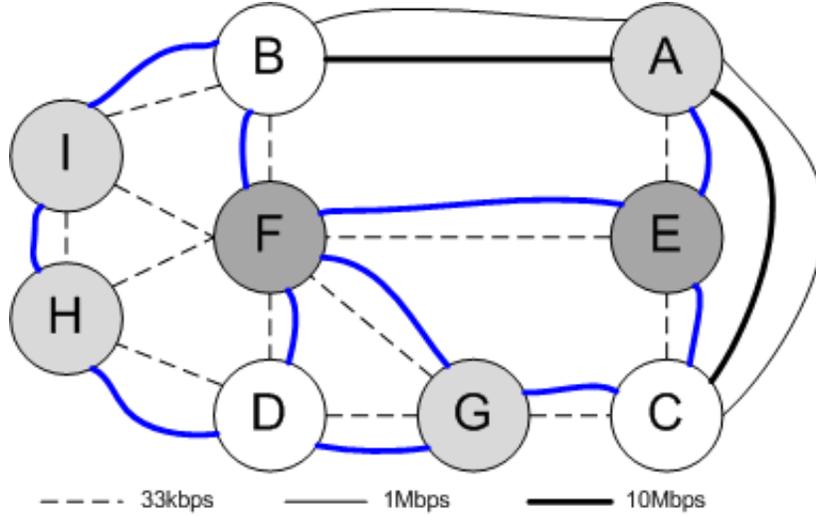


Figure 4.4: Simplified network provisioning graph with sources (light gray), sinks (white), and routers (dark gray) identified. Also includes provisional links in blue.

ogy because the cost of the edges is a sub-additive monotonic function, meaning that the unit cost of bandwidth decreases as demand, and consequently allocated supply, increases.

When a single function defining this cost scheme is applied to all edges, the topology is considered to be *uniform*. On the other hand, *non-uniform* networks define a different cost function for each pair of nodes. This supports more complex topologies that must account for pre-installed infrastructure components or other complicating factors. Furthermore, the specific algorithm we use extends this concept by representing the non-uniform piecewise functions as multiple edges between the appropriate pairs of nodes. A fixed cost c_e and an incremental cost l_e (per unit bandwidth transferred over edge) are assigned to each edge e . The total cost of a particular solution can be obtained by choosing a set of edges E' and, for each source/sink pair (s_i, t_i) , a path from s_i to t_i using edges in E' that minimizes the value

$$c(E') + \sum_{i=1}^h d_i \cdot l_{E'}(s_i, t_i),$$

where $l_{E'}(s_i, t_i)$ is the total incremental cost of the chosen path between u and v . Our aim is to show how to exploit the following theorem [CHKS06] for ERN readiness assessment.

Theorem 1 *There is a polynomial time algorithm for the multi-commodity buy-at-bulk network provisioning problem with an $O(\log^4 h)$ approximation ratio, where h is the number of source/sink pairs.*

The algorithm to which this theorem refers is not ideally suited to this task, as we will show later, but it should motivate further research into improved algorithms that are better matched with the specific requirements of this problem. Regardless, to use the algorithm we need to supplement the graph representation of our problem with the cost information needed by the multi-commodity buy-at-bulk algorithm. First, we must define pairs of source and sink nodes. In fact, the designation of source or sink is relatively unimportant in the case of symmetric full-duplex networks, like the ones we are primarily concerned about, since if it is possible to route traffic in one direction on a link, it is possible to route an equal amount of traffic in the other direction. However, this may not be true for all links such as DSL connections, and delay-tolerant networks that use mobile objects such as buses to transport data.

For asymmetric links like DSL we simply use the smaller bandwidth as the link's total bandwidth. Truly unidirectional links are fundamentally incompatible with the algorithm in [CHKS06]. Thus, we require that unidirectional routes be accounted for manually. If a pair of nodes must communicate over unidirectional routes, or it is optimal for them to do so, they should be removed from the problem formulation during automated analysis.

Now, it is necessary to map the edges in our initial problem graph to edges with a scalar cost value suitable for analysis using this algorithm. The primary difficulty to be overcome in this task is to ensure that the algorithm does not exceed the bandwidth limitations of individual links. There is no notion of bandwidth limitation in the algorithm, so this must be accomplished indirectly by engineering the fixed and incremental costs of the edges. For example, a model may consider both pre-installed advanced meter mesh connections and supplementary dedicated ERN connections. The meter connections have low bandwidth but are pre-installed, so they have a relatively low fixed cost, whereas dedicated networks have higher bandwidth but also have a higher fixed cost, since they must be installed specifically for emergency-response purposes. Thus, if metering networks provide sufficient bandwidth to support the required communications, they should be used. However, if additional bandwidth is required, the algorithm

should select dedicated networking components.

This can be accomplished by ensuring that the artificial cost of the low-bandwidth network exceeds the cost of the high-bandwidth network before the bandwidth demand exceeds the supply provided by the low-bandwidth network. One important point to note in this discussion is that whenever two nodes are connected exclusively by edges with bandwidth lower than the total network demand, that pair of nodes must also be connected by a provisional edge that has practically infinite bandwidth from the standpoint of the provisioning algorithm, which can be accomplished by making its bandwidth greater than the total network demand. The cost of such an edge may be prohibitive for real deployments, but it must be included to ensure that the algorithm does not assign more demand than the low-bandwidth edge can handle. In an adequately-connected network, the expensive infinite link should never be selected in an approximately optimum solution unless it is absolutely necessary. Due to the linear incremental cost function required by the algorithm, our approach artificially inflates the incremental costs of low-bandwidth links, even at relatively low utilization levels. A quadratic or exponential cost function would provide better results. Addressing this issue is an important research problem.

Other considerations besides bandwidth are very important in this context. Thus, we must factor the attributes described earlier into the cost determination. The first attribute, **cost**, is included directly as the fixed cost of the edge. The second attribute, **reliability**, is a non-monetary attribute that should be integrated in such a way as to cause the algorithm to prefer edges with high reliability over those with low reliability. To do this, we decrease the fixed cost of high-reliability edges relative to low-reliability edges. High **latency** is generally undesirable in most emergency-response applications, so low-latency links should be preferred over high-latency links. However, the relative importance of latency as compared to reliability and other factors varies depending on the expected application usage on the network. Generally, high-latency links are not well-suited to real-time voice communications or even interactive text messaging, so latency is important in typical ERN usage scenarios.

To make this discussion more precise, we define several equations to determine the artificial fixed cost of an edge:

Definition The fixed cost of an edge e is expressed as $c_e = (\text{equipcost}(e) + lw \cdot$

$latency(e)/reliability(e)$, where $equipcost(e)$ is the cost to install and maintain e , lw is the weight accorded to latency in this network, $latency(e)$ is the expected latency of e , and the overall reliability of e considering all disasters is expressed as $reliability(e) = \sum_{d \in \mathcal{D}} (p(e, d) \cdot dependability(e, d))$, where \mathcal{D} is the set of disasters that may occur, $p(e, d)$ is the probability that a particular disaster d will occur, and the dependability of an edge e when exposed to disaster d is expressed as:

$$dependability(e, d) = \sum_{h \in \mathcal{H}_d} \left(\frac{p(e, h) \cdot availability(e, h)}{susceptibility(e, h)} \right),$$

where \mathcal{H}_d is the set of all hazards that may occur in disaster d , $p(e, h)$ is the probability that a particular hazard h will occur during the disaster, $susceptibility(e, h)$ is the probability that h will degrade or destroy edge e , and $availability(e, h)$ is the probability that e will be made available for emergency communications when exposed to h .

Next, we define the artificial incremental cost of an edge:

Definition The incremental cost of an edge e is expressed as:

$$l_e = \max \left(l_{nextbigger(e)}, \left(bwchrg(e) + \max \left(0, \left(\frac{c_{nextbigger(e)} - c_e}{capacity(e)} - bwchrg(e) \right) \right) \right) \right),$$

where $nextbigger(e)$ is the edge that is parallel to e and has the next lowest bandwidth, $capacity(e)$ is the bandwidth of e , and $bwchrg(e)$ is the monetary cost of transmitting an additional unit of data along e .

This definition ensures that the capacity of an edge is never exceeded, because to do so would cost more than to select the next edge with adequate capacity. Of course, these costs do not necessarily correspond to monetary costs, but they do factor in the appropriate monetary costs while also integrating the other considerations that are critical in ERNs.

The graph that results from designating pairs of nodes as sources and sinks, adding effectively infinite bandwidth links where necessary, and assigning edge costs is suitable for analysis using [CHKS06]. The algorithm results in an approximately optimal subgraph that provides adequate bandwidth to satisfy all demand pairs.

Conjecture 1 *The result of the algorithm is a feasible and approximately optimum network that satisfies the expected bandwidth requirements of all pairs of nodes.*

In many instances, the reliability provided by the resulting network may be inadequate, since the network may be interrupted by a single break in one of the links. Thus, it may be necessary to iterate the algorithm, to develop redundant networks. By removing the edges in the approximately optimum subgraph from the original graph, recalculating the edge costs to ensure that bandwidth limitations are not exceeded, and re-running the algorithm on the new graph, a fully-redundant network infrastructure can be developed. This process can be repeated as many times as is economically feasible to develop a network with a corresponding level of redundancy.

The preceding workflow is only capable of selecting among provisional infrastructure enhancements that are inserted into the network graph by whatever entity is performing the workflow. Some of these enhancements may be in obvious locations, such as at the tops of telephone poles or municipal buildings, where they are relatively easy to install. However, these locations may not be sufficient to provide complete network connectivity. Simple algorithms could be developed to compute the geographical coverage of existing actual and provisional network assets, and then highlight geographical regions that are not within this coverage area and are likely to require network connectivity. Then, the individual performing the emergency-response provisioning workflow could determine feasible locations for provisional infrastructure enhancements within the highlighted regions.

4.3.2 Example Scenario

Let us return to our previous example of a municipal intranet deployed in parallel with an AMI network to see how this approach could be helpful and to point out remaining research questions. Assume the intranet uses a standard IP network based on Ethernet technologies from end-to-end, the AMI network uses 802.15.4 running an IP layer, and the cameras use 802.11b with IP. Furthermore, assume that some disaster divides the intranet into several disconnected fragments. Then, to support communications, those fragments must be reconnected to each other. This can occur if the secondary networks route packets between the disconnected fragments. Obviously, this requires an interface between each Ethernet, 802.11b and 802.15.4 fragment. 802.11b and 802.15.4

radio interfaces are very inexpensive, so this is not a significant problem. In fact, the gateway nodes could be deployed after the disaster occurs, although this will introduce potentially significant latency into the recovery process. To determine the locations where gateway nodes will most likely be needed, whether they are pre-installed or not, all likely gateway node locations can be included in the graph provided to the algorithm. The algorithm will select those nodes that are most likely to provide critical support to network connectivity. If the nodes are then pre-installed, they will provide constant ERN support. Otherwise, emergency responders will know precisely where network enhancements should be installed to provide maximum effect, in contrast with current ad-hoc approaches of emergency response infrastructure deployment.

Of course, even if connectivity is supported, the AMI network must be configured to permit packets from the municipal intranet to be routed. If we assume that IP is in use on all networks we can be assured that all nodes will be reachable in a connected graph. However, the bandwidth demands between source and sink nodes may not be satisfied if IP selects different routes than those chosen by the network design algorithm, since IP routing algorithms are best suited for hierarchical network organizations. A hierarchy is unlikely to exist in ERNs, so we believe that routing algorithms more akin to those used in ad-hoc networks and ZigBee may be more appropriate. Thus, we focus on such algorithms in this dissertation.

5 ERN Protocol Specification

5.1 Network and Transport Layers

The fundamental philosophy of our system is to layer ad-hoc network protocols on heterogeneous transport protocols. In this section, we describe the frame formats that correspond to the high-level network layer, layer 1 in Figure 5.1.

This layer has two frame formats, one each for unicast and broadcast messages, depicted in Figure 5.2. All messages specify a source ERN address, a broadcast flag that is set only for broadcast messages, and a message ID that serves to distinguish messages sent by a single host.

Broadcast messages specify a TTL value that limits their broadcast radius, a set of flags, and a payload that always contains a transport frame. Routers must place upper bounds on TTL values to prevent network abuse. Currently, no flags are defined for broadcast messages.

Unicast messages specify a destination ERN address, a set of flags, and a transport frame payload. Currently, only the lowest bit of the flags field is used, to indicate whether the message is encrypted and authenticated. Optionally, a cryptographic hash can be included if QoS controls are in use. The hash is a member of a hash chain that is

Low-Level Application Layer	ERN Application Layer (Routing, Text, Voice, etc.) [3]
	ERN Transport Layer (Data Messages with App. Info) [2]
	ERN Network Layer (Addressing, Unicast, Broadcast) [1]
Low-Level Transport Layers (UDP, ZigBee Transport)	
Low-Level Network Layers (IP, ZigBee Network)	
Low-Level Data Link and Physical Layers (Ethernet, 802.11, 802.15.4)	

Figure 5.1: ERN network layers and their relationship to the layers of the underlying networks used to form the ERN.

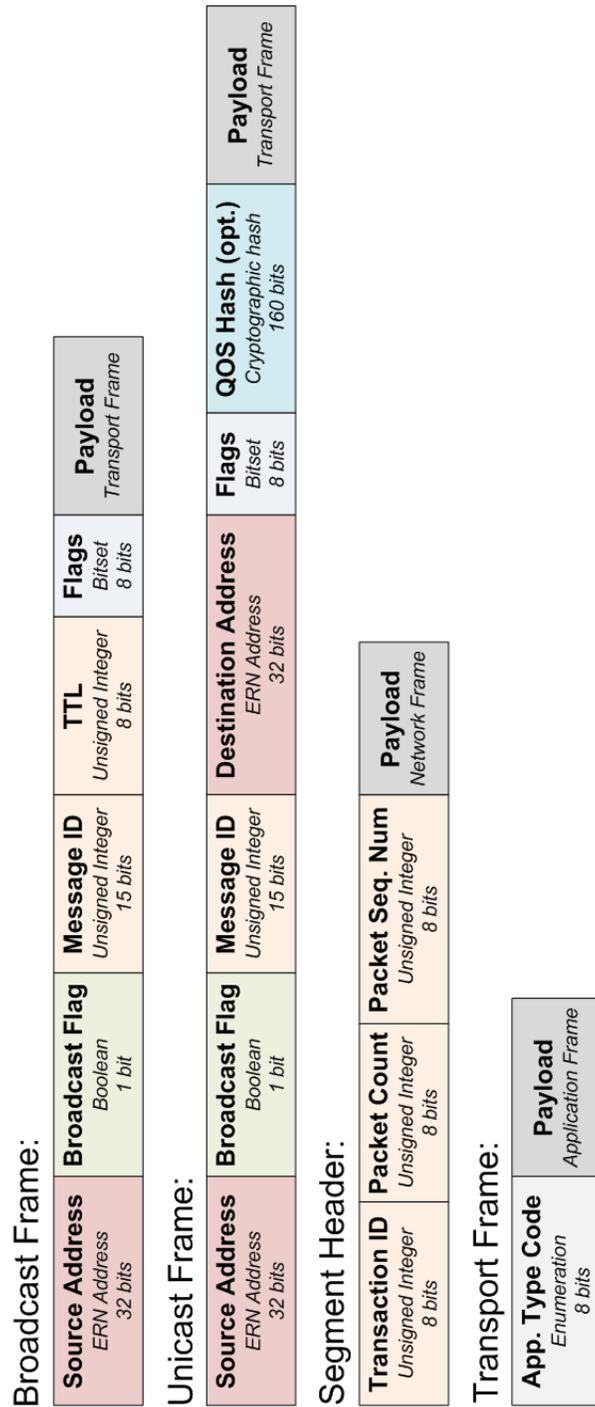


Figure 5.2: ERN network and transport layer frames, plus the segment header.

used to convince routers that the message is entitled to a specific QoS level. This will be more thoroughly explained later.

The transport header contains a field that determines what type of application frame follows the header. A unique code has been allocated for every type of application frame that will be discussed in the next section.

All messages are segmented before being transmitted on an underlying network, since the underlying networks that we target are packet-oriented and often have MTU sizes that are too small to accommodate complete ERN messages. Packets are re-assembled and then re-segmented at every hop in the ERN, to optimize efficiency on routes that include underlying networks with different MTU sizes. The segment header contains a transaction ID that is used to distinguish separate messages originating from the same underlying network host, the total number of segments in the current message, and the sequence number of the current segment.

5.2 Application Layer

5.2.1 Routing

The AODV routing algorithm is used to route ERN messages [PR99]. The frame formats used to implement AODV are shown in Figure 5.3. Every route request specifies the destination ERN address of the node to which a route is being sought, the sequence number of the latest routing request or response that the issuing host is aware of, the endpoint ID of the node issuing the request, and the current routing sequence number of the host issuing the request. The sequence numbers are used to prevent older routing information from supplanting fresher routes. When a node receives a route request for an address that exists within its routing table, it checks the sequence numbers of the stored route and the request, and if the stored route is fresher, sends a response to the requester back along the same underlying network path that was used to propagate the request. The 128-bit endpoint IDs are used to distinguish between different nodes that have accidentally selected identical ERN addresses. This provides some of the benefits of a very large address space, by making the probability of endpoint ID collisions very small, while also providing the efficiency advantages of a small address space for most messages. Of course, if the node being sought itself receives a request, it will

Acknowledgment (Unicast):

Ack'ed Message ID <i>Unsigned Integer</i> 8 bits	Ack'ed Message Type <i>Enumeration</i> 8 bits	Ack'ed Message Destination <i>ERN Address</i> 32 bits
---	--	--

Route Request (Broadcast):

Dest. Addr. <i>ERN Address</i> 32 bits	Dest. Seq. Num. <i>Unsigned Integer</i> 16 bits	Endpoint ID <i>Unsigned Integer</i> 128 bits	Source Seq. Num. <i>Unsigned Integer</i> 16 bits
---	--	---	---

Route Response (Unicast):

Dest. Addr. <i>ERN Address</i> 32 bits	Dest. Seq. Num. <i>Unsigned Integer</i> 16 bits	Endpoint ID <i>Unsigned Integer</i> 128 bits	Generation Time <i>Unsigned Integer</i> 64 bits
---	--	---	--

Duplicate Address Notification (Unicast):

(null)

Session Establishment Request (Unicast):

Nonce <i>Binary Data</i> 160 bits	Request Time <i>Unsigned Integer</i> 64 bits	Requester Cert. <i>ECC Certificate</i> > 576 bits	Message Signature <i>ECC Signature</i> 388 bits
--	---	--	--

Figure 5.3: Acknowledgment and routing frame formats.

unconditionally send back a response with the latest sequence number. Otherwise, the request is re-broadcast. Each response is a unicast message that contains the destination address originally requested, since the response message may originate at a different node. It also contains the sequence number associated with the route, the endpoint ID of the respondent, and the time when the response was generated.

To acquire an ERN address, each node must randomly generate an address and then attempt to determine whether that address is already in use. To do this, each node must issue a routing request for that address. Other nodes that receive the request will handle it specially, since its source and destination addresses match. If any node that receives the request has routing information for the address in question, it will send a duplicate address notification (described below) to the new node that is attempting to use the address, without updating its local routing tables. The new node simply repeats this process until an available address has been discovered. Of course, this process does not guarantee that the new address is actually available, since another node in a disconnected segment of the ERN may be using it, or the request may not have reached

a node with knowledge of the address, even if such a node exists in the local network segment, due to transient conditions. Thus, an additional mechanism is available to help resolve such conflicts.

The duplicate address notification message contains no additional information in its payload, since it unequivocally indicates that its destination is using a duplicated address. It is generated whenever a node detects that there are duplicated addresses present on the network, which can occur at many points in the routing process. For example, that situation can be detected when a node requests a route to the duplicated address, and receives more than one response with different endpoint nonces. The node that requested the route generates a notification for each node using the duplicated address, and sends them along the appropriate routes. When a node receives such a notification, it must immediately request a new address, as described previously. It is highly unlikely that endpoint IDs will randomly collide, so this is a robust mechanism.

Some unicast messages are confidential and must be encrypted before being transmitted. To accomplish this, a session key can be established between a pair of nodes upon demand. To support this functionality, we developed a simple key exchange protocol inspired by JFK [ABB⁺04]. Anonymity is unimportant in an ERN, so we did not actually implement the full protocol. Likewise, JFK provides confirmation to both parties at the end of the exchange that a key has been successfully established. To eliminate the extra messages required for that confirmation, we simply assume that the key exchange completed successfully, and indicate failure if subsequent decryption attempts fail. Since our exact protocol has not been carefully analyzed, it may have security vulnerabilities, so a full implementation of JFK or some other carefully analyzed protocol must be used in production environments.

To request that a session key be established, the initiator sends a session establishment request to the other node as a unicast message. This request contains a nonce, to ensure that every session has a unique key, a timestamp to prevent replay attacks, the requester's certificate, and a signature over the whole message. The signature proves that the initiator actually has use of the private key corresponding to the provided certificate, and has requested that a session be established with the provided nonce soon after the indicated time. The replay prevention threshold must be high enough to accommodate high-latency ERN links and loosely-synchronized clocks.

Upon receiving a session establishment request, if the recipient wishes to establish a session with the sender, the recipient must perform an Elliptic-Curve Diffie Hellman (ECDH) key exchange using the sender certificate and nonce provided in the request, and the local certificate. Then, the recipient must send back a similar request to the sender, including the same nonce and the original recipient's certificate. The original sender then repeats the key derivation process, and after both requests have been acknowledged, a session is assumed to have been successfully established. Then, whenever a confidential message is sent between the two parties, the session key is used to encrypt and authenticate it with the Counter with CBC-MAC (CCM) block cipher mode for AES-128. Only the body of the transport payload is encrypted (not the type code), but the entire message is authenticated (excluding the segment headers, of course).

Obviously, the key establishment process requires at least one roundtrip before communications can be established, which makes it unsuitable for most Delay-Tolerant Networks (DTNs) [HABR05]. Thus, it is likely that all messages sent over DTNs will be sent in cleartext.

5.2.2 ERN Applications

A variety of high-level application frames can also be inserted into transport frames. Some of these frames are depicted in Figure 5.4. We describe each of these frame types in this section.

Image frames contain data to construct a visual image. Standard formats such as JPEG, GIF, and PNG must be supported. Every image comprises a descriptive caption, an image type specification (in case the image format does not contain a magic number), and binary image data.

Audio frames contain streaming audio data packets. This type of frame is intended to transmit voice data, and the Speex codec is very well suited to encoding voice data [Val06]. Thus, each audio frame must contain a sequence of Speex-encoded packets. Each packet must be preceded by a 16-bit length specifier. The number of packets can be variable and even automatically adjusted for optimal performance.

Text messages simply contain an arbitrary string of Unicode characters.

Panic button requests specify the location of the individual requesting assistance,

if known, and also include a free-form comment from the requester. Panic button responses are formatted similarly, but specify the respondent location and include a comment from the respondent. Each response also specifies the rescuer that originally generated the response. The reason for this will be clear shortly. Additionally, each response includes a response type field. That field specifies whether: *a)* the respondent is providing a conditional offer of assistance, pending responses from other possible rescuers, *b)* the respondent is committed to performing the rescue, or *c)* the respondent is rescinding a previous offer or commitment. To coordinate a response effort, several actions must take place. First, the victim must broadcast a properly-formatted panic button request. Any rescuer in range that is willing to assist must respond with a panic button response. Upon receiving each request, the victim must forward every panic button response that has already been received to the new respondent. Then, the victim must forward the new response to all previous respondents. After waiting for an arbitrary period of time, depending upon the urgency of the request and other factors, some of the respondents should respond with a firm commitment, perhaps after negotiating with the other respondents.

Emergency indications are used by nodes to inform other nodes that they have detected an emergency. Each node uses these indications to determine when it should provide ERN services. Each indication has a bitset of type flags, used to indicate what types of emergency conditions have been detected by the originating node. For each bit that is set, the indication includes an 8-bit signed integer that specifies how confident the indicating node is that such an emergency condition actually exists. The confidence level is signed, because a node is permitted to issue a negative indication if it has observed positive indications from other nodes and possesses evidence suggesting that the specified emergency conditions are in fact not present. Every indication must also include time of occurrence and duration specifications. Finally, all indications must be signed, and the certificate of the originator must be included in the indication. Emergency indications are broadcast to all nodes within a specific broadcast radius. Thus, the indication may reach nodes that are not affected by the emergency conditions, but this is actually useful, because those nodes are probably well-positioned to provide ERN services for the adjacent nodes that are affected by the emergency.

ERNs must also provide a service for locating nodes according to the human-

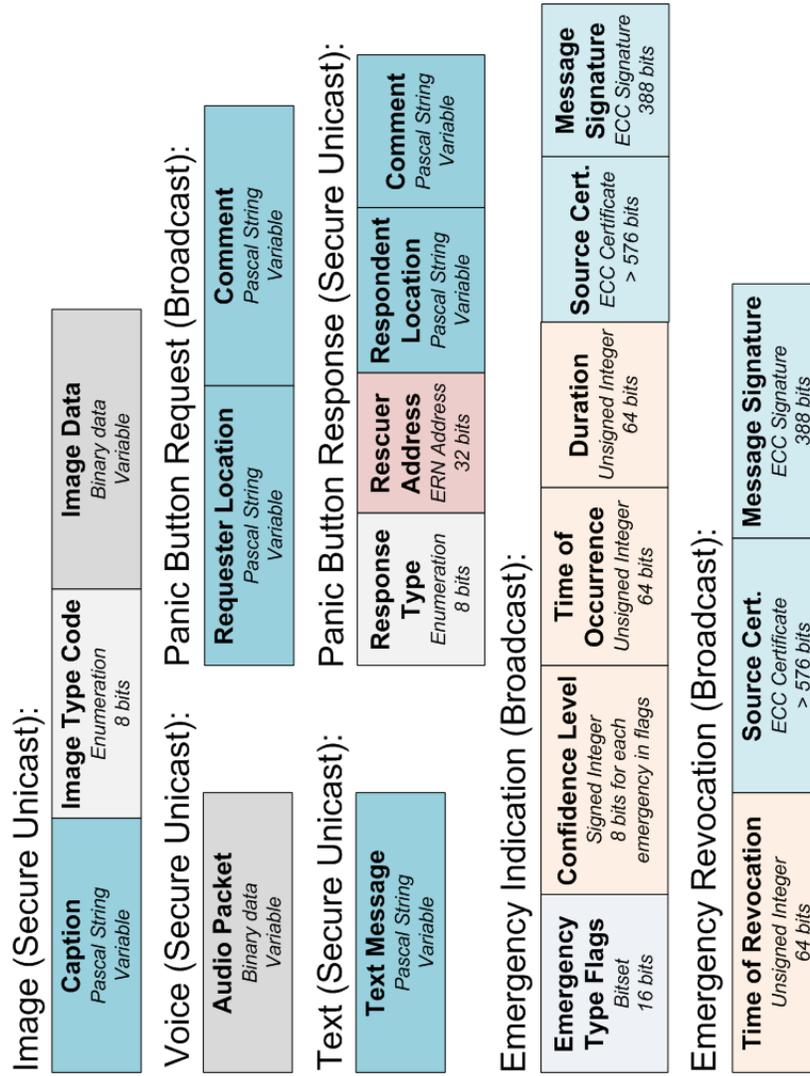


Figure 5.4: High-level application ERN frames.

recognizable attributes of their associated entity, most likely a human using the device. To support this, a few nodes on each network must provide simple database functionality. Each database has a mapping from a legal name to the address of the node used by the entity with that legal name. Such databases are most likely to be operated by emergency response organizations. Each database must periodically broadcast a presence announcement, which includes a signed copy of the database node certificate, to allow nodes to recognize databases operated by reputable organizations and avoid any malicious databases that may arise.

Whenever a node receives a database announcement, it must first ensure that the announcement is from a trusted node. If so, it must submit a registration message containing the legal name associated with the node's user, and any comments that the user wishes to include, such as information on current health status or plans for escaping the disaster. When a node wishes to locate another node associated with a particular individual, it must submit a lookup request to any trusted databases that it is aware of. Each request contains the desired legal name. The database will respond with the original registration message submitted by the desired individual if it is available.

Often, a victim wishes to locate another node that is associated with a person with a specific role, rather than a specific identity. This is also supported, and does not require any interaction with database nodes. Instead, the node performing the search broadcasts a request that specifies the desired role, and a request ID used by respondents. Any node within range that possesses the desired role must respond to the request with the request ID. Then, the requester is free to establish a session with any respondent using text, image, and/or audio messaging.

5.2.3 QoS

Optionally, an ERN can implement Quality-of-Service (QoS) controls, so that rescuers and other authorized nodes can have greater priority than other nodes. QoS can only be guaranteed uni-directionally between a specific pair of nodes, along a specific route. To request guaranteed QoS, one of the nodes must issue a QoS request, which specifies the time at which the guarantees should be initiated, the amount of time for which they will be required, the amount of bandwidth that is desired, the maximum desired latency, a cryptographic hash chain seed that is used to inexpensively authenticate the

Role Lookup Request (Broadcast):

Request ID <i>Unsigned Integer</i> 32 bits	Role Specification <i>Enumeration</i> 16 bits
---	--

Role Lookup Response (Unicast):

Request ID <i>Unsigned Integer</i> 32 bits

Node Info DB Announcement (Broadcast):

Database Cert. <i>ECC Certificate</i> > 576 bits	Message Signature <i>ECC Signature</i> 388 bits
---	--

Node Info DB Registration (Unicast):

Legal Name <i>Pascal String</i> Variable	Notes <i>Pascal String</i> Variable
---	--

Node Info DB Lookup Request (Unicast):

Legal Name <i>Pascal String</i> Variable

Node Info DB Lookup Response (Unicast):

Matching Node Info DB Registration <i>Encapsulated Payload</i>
--

Figure 5.5: Application-level ERN frames for node information lookups.

QoS Request (Unicast):

Time of Request <i>Unsigned Integer</i> 64 bits	Duration <i>Unsigned Integer</i> 64 bits	Desired Bandwidth (KB/s) <i>Unsigned Integer</i> 32 bits	Desired Latency (ms) <i>Unsigned Integer</i> 32 bits	Hash Chain Seed <i>Cryptographic Hash</i> 160 bits	Source Cert. <i>ECC Certificate</i> > 576 bits	Message Signature <i>ECC Signature</i> 388 bits
--	---	---	---	---	---	--

QoS Response (Unicast):

Start Time <i>Unsigned Integer</i> 64 bits	Duration <i>Unsigned Integer</i> 64 bits	Guaranteed Bandwidth (KB/s) <i>Unsigned Integer</i> 32 bits	Guaranteed Latency (ms) <i>Unsigned Integer</i> 32 bits
---	---	--	--

Figure 5.6: Quality-of-Service control frames.

node once QoS guarantees have been established, the certificate of the initiator, and a signature over the entire request. The initiator sends this request as a unicast message to the destination node, and each router along the path must respond to the request if it is willing and able to provide the requested QoS guarantees. In fact, the last router in the path must generate the first response, and each router along the path back to the initiator must either adjust the response to reflect any limitations that are in place, and then continue forwarding it. Ultimately, the initiator will either receive a response that specifies the QoS guarantees provided along the entire route, or no response if no guarantees can be provided or QoS is not supported. If the guarantees are inferior to those requested by the initiator, it must decide whether they are acceptable. If so, it must issue a new request that contains the parameters specified in the response, so that any downstream nodes that have reserved more resources can adjust their reservations appropriately. Otherwise, it must issue a new request with the most permissive requirements possible (zero bandwidth and maximum latency), so that the QoS reservations can be canceled. If no QoS response is received, a similar cancellation must be issued, in case a downstream node reserved resources.

For every unicast message that requires the QoS guarantees, the QoS flag must be set, and the next hash value in the hash chain must be inserted after the flags field. Once the hash chain has been exhausted, a new QoS negotiation must occur.

6 Implementation and Testbed

To evaluate the ERN architecture that we have proposed, we implemented a prototype that is capable of retasking ZigBee and IP network devices to provide ERN services. We evaluate the basic functionality of the system, and then evaluate its performance for different usage scenarios.

6.1 Functional Prototype

6.1.1 Overview

To demonstrate the feasibility of our system, we have developed a functional prototype. All nodes run a Java application that implements all of the protocols described above, with the exception of QoS, which is only partially implemented and not tested in our experiments. The application has a graphical front-end that is used to control its operation and monitor the state of the ERN subsystem of the node. It has selectable backends for controlling and using Maxstream XBee ZigBee radios and standard IP networks such as Ethernet and Wi-Fi. Design patterns were used extensively throughout the entire application, making it easily extensible. In the remainder of this section, we describe the major features of this prototype software, after we introduce the ZigBee radio hardware that it controls.

Maxstream XBee Series 2 ZigBee radios are embedded modules that each contain a microcontroller, ZigBee chipset, and antenna. Multiple antenna types are supported, and the choice of antenna can significantly affect the range of each node. The theoretical maximum data rate between two nodes is 250 kbps. The maximum indoor range between two nodes is claimed to be 40 m, but no antenna type is specified. Outdoors, the nodes should be able to ideally attain a range of 120 m. The maximum transmit power is 2 mW, and the receivers have a maximum sensitivity of -98 dBm. Drop-in compatible replacement radios are available that have increased output power and re-

Parameter Name	Description	Register Value
AP	Modem command mode	UNESCAPED
BD	UART data rate (baud)	BD115200
CH	Channel (XBee: 0x0B-0x1A, PRO: 0x0C-0x17)	23 (8'h17)
DO	Digital I/O port 0 configuration	ASSOC_RTS_CTS_IND
HV	Hardware version	16'h1941
ID	Personal-area-network ID	16'h0777
SH	High serial number	1286656 (32'h0013A200)
SL	Low serial number	1074401530 (32'h400A10FA)
MY	16-bit source address	16'h4B4E
NI	Node identifier	USB-WHIP-1
NT	Node discovery timeout	60 (8'h3C)
PL	Power-level (0-4, see docs)	4 (8'h04)
SC	Channel list for energy and activity scans (bit flags: 0x1=11 to 0x4000=26)	16'b00011111111111110
SD	Scan duration (XBee 1: 0-15) (XBee 2: 0-7)	3 (8'h03)
SM	Sleep mode	DISABLED_12
SP	Cyclic sleep period (XBee 1: 1-0x68B0 ms, XBee 2: 0x20-0xAFO x 10 ms)	32 (16'h0020)
ST	Sleep timeout	5000 (16'h1388)
VR	Firmware version	16'h1341
NJ	Node join time (0-0x40 sec, 0xFF for always)	255 (8'hFF)
AO	RX indicator mode	EXPLICIT_RX_DATA_IND
JN	Join notification mode	DISABLED
PM	Power boost enabled	true
SN	Maximum number of consecutive sleep periods	1 (8'h01)

Figure 6.1: Configuration utility for setting parameters of Maxstream XBee Series 2 ZigBee modems.

ceive sensitivity, and can attain an outdoor range of 1.6 km. Each node is connected to a serial port, and development boards are available that include a USB-to-serial bridge chip. Multiple firmware versions are available for these nodes, and can support different modes of operation. By default, the nodes operate as serial-port extenders and simply send all data that enters the serial port to another designated node, and send all radio data that is designated for itself through the serial port. The only exception to this is a set of AT modem commands that are specially interpreted by the device and used to query and update configuration settings. Alternatively, an “API-mode” firmware can be used that can accept and produce compact binary packets that closely correspond to the actual packets transmitted in the ZigBee network. AT configuration commands are also supported by this firmware version, but are implemented as special binary packets. We use API-mode firmware in all of our experiments.

To control the serial-connected modems, we used the RXTX cross-platform serial port library for Java. All modems were configured to communicate with the PC at 115200 bps. To set configuration parameters on the modems, we implemented the AT modem commands in our Java interface, as shown in Figure 6.1. We configured all modems to use channel 23, which is the one least likely to conflict with Wi-Fi.

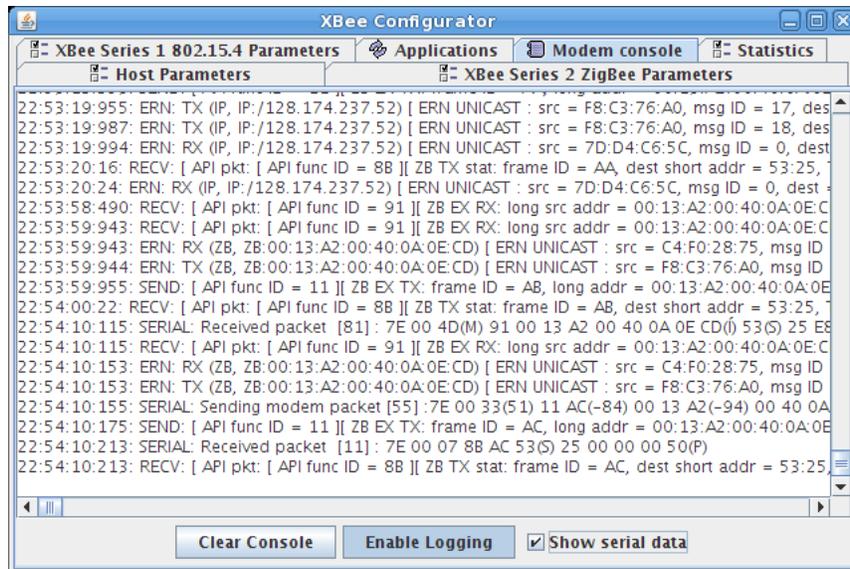


Figure 6.2: Network audit console.

To monitor the operation of the system, we implemented a logging console that outputs information on all messages sent and received by the ERN layer as well as any underlying network layers (ZigBee and IP in the current prototype), as shown in Figure 6.2. Even serial port data can be displayed if required. Unfortunately, it is very computationally expensive to update the log, so it must be disabled during certain transactions, such as audio messaging.

We slightly modified the network packet formats specified for ERNs to include a timestamp, so that we can perform network latency and throughput measurements. These timestamps add a 64-bit overhead to every packet.

6.1.2 Prototype ERN Applications

Several applications are implemented in our Java framework, and the ERN application is one of these. After they select the ERN application, the user is presented with the main panel shown in Figure 6.3. The node's ERN address is shown at the top of the panel, beside a button used to acquire a new ERN address on demand. If a ZigBee modem is connected, its 64-bit 802.15.4 address and 16-bit local network address are displayed. If IP networking is enabled using the button labeled to that effect, the IP address of the node is also displayed. That is followed by a text area describing the certificate associated with the node. The next segment of the panel indicates whether

ERN services have been activated, followed by a measure of the minimum indication entropy among the various types of disasters, the entropy threshold used to determine whether ERN services can be activated, followed by individual confidence indications for each type of recognized emergency. Any category of emergency with entropy above the entropy threshold has a confidence level of approximately 0% (discretization errors are introduced by the conversion from an 8-bit confidence level to a confidence percentage). These indications are followed by a list of other nodes that have been detected on the network. Each node has a list entry displaying its ERN address, the freshness of the latest route to that address (sequence number), the number of hops required to reach the specified node (or 20 if unknown), the native address of the next router along the current route to that node, and the certificate of the node if available. The certificate becomes available when a session is established with the node. This is also indicated with a “[SESS]” tag at the beginning of the node’s entry. When a node in the list is selected, its certificate is displayed in the text area beneath the list. The “Create Comm Session” button is also enabled. Pressing that button presents a list of message types that can be transmitted from the node: text, image, or audio. The “Perform Unassociated Action” button presents a list of operations that can be performed without selecting any particular node. A node information database can be created on the local node, a panic button request can be issued, etc. We go over these actions in detail in the following paragraphs.

Seven distinct actions can be performed without being directed at any particular node on the network. The first three concern node information databases. First, one or more nodes must create node information databases and announce their presence on the network. On the node that creates the database, a list of registered nodes appears on the graphical testbench, along with a button that can be used to initiate a re-announcement of the database’s presence (Figure 6.4). All nodes that receive a database announcement show a dialog that requests their users to register with one or more of the new databases (Figure 6.5). Finally, any node can query a database for node information. First, the query must be entered in a simple text dialog, then a search dialog is created (Figure 6.6). The querying node must select the database to search in that dialog and then initiate the search. If a response arrives, the dialog is updated with the results of the search, and the user is permitted to optionally establish a session key with

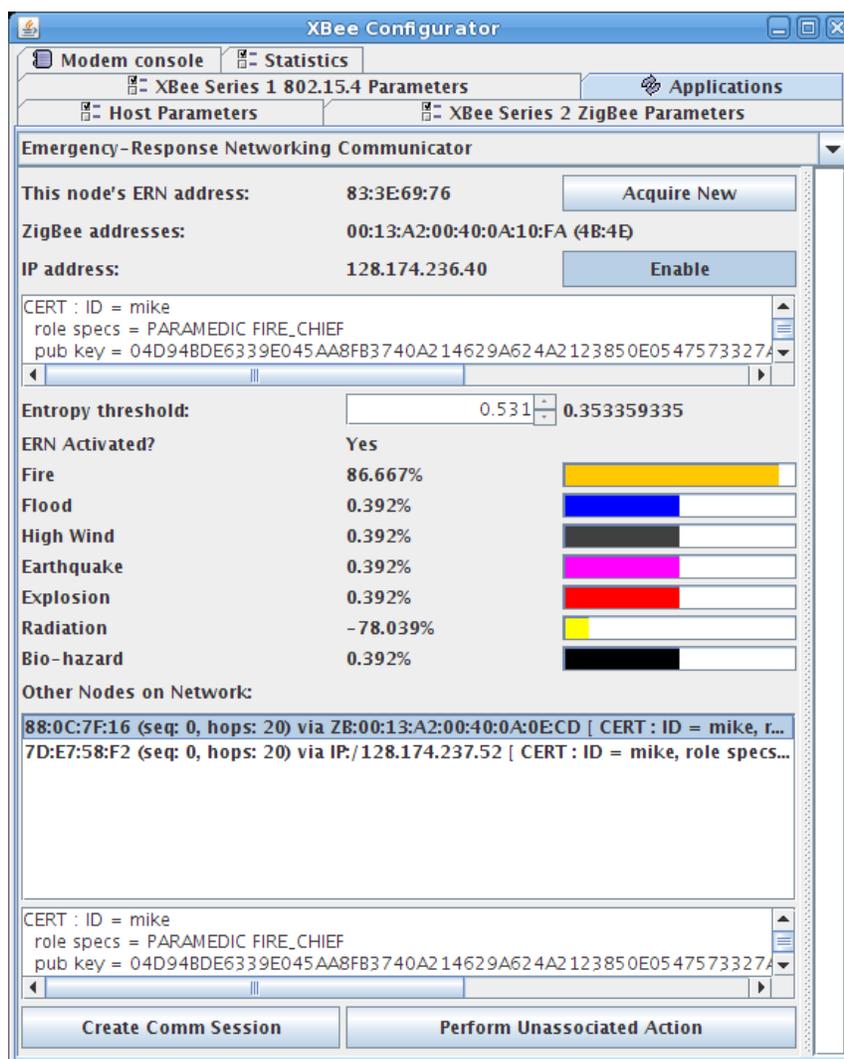


Figure 6.3: Main ERN application panel.

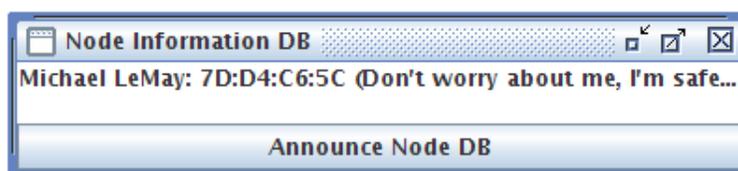


Figure 6.4: Node information database management dialog.

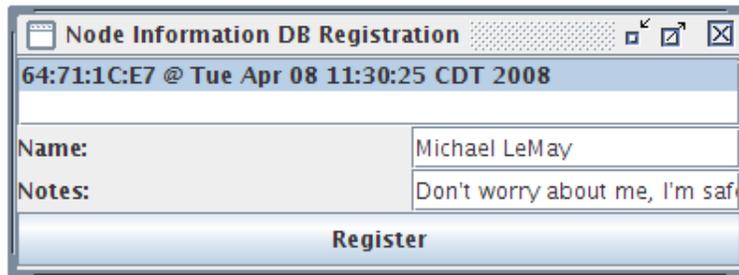


Figure 6.5: Node information database registration dialog.

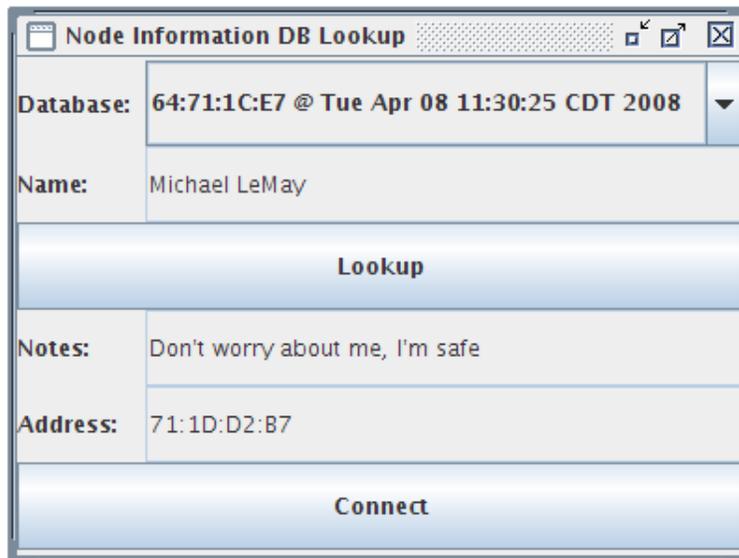


Figure 6.6: Node information database query dialog.

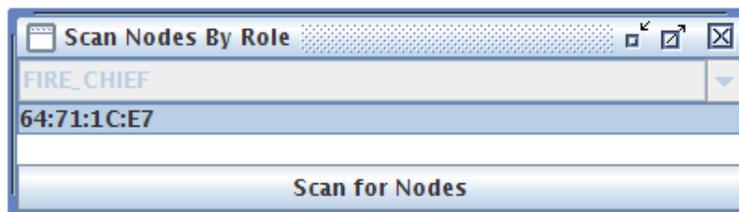


Figure 6.7: Dialog to search for nodes by role.

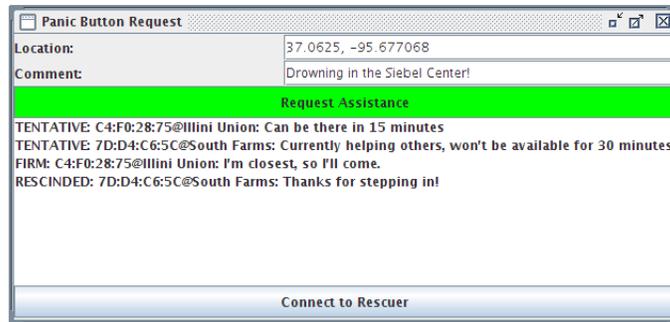


Figure 6.8: Panic button request dialog.

the discovered node, in preparation for future communications. Even in the absence of any semantic node databases, it is still possible to search for other nodes that possess specific roles. A search dialog is created that simply contains a drop-down list of all possible roles and a list of respondents that possess that role and are within the broadcast radius of the node that initiated the search (Figure 6.7).

Panic button requests can be transmitted from a dialog that permits the victim to describe their current location, enter a comment indicating the assistance that is required, and then issue the request (Figure 6.8). The dialog contains a list of all responses that are received. The color of the button used to issue the request indicates the current status of the request. If it is yellow, no request has been issued yet, or only tentative responses have been received. If it is red, a request has been issued, but no responses have been received, or all respondents have rescinded their offers of assistance. Finally, if it is green, at least one respondent has firmly committed to providing assistance. The panic button response dialog is similar, but requires the respondent to select which type of response should be transmitted.

Our current implementation does not use inputs from any actual sensors to detect emergency conditions, so we manually transmit arbitrary emergency indications instead. The emergency indication dialog has a spinner for each type of emergency condition that is used to indicate the confidence the originating node has that such emergency conditions are in effect (Figure 6.9). After these confidence levels have been adjusted, the originator must select an expiration time for the indication, and then issue the indication message.

When a semantically-rich emergency announcement must be propagated to a large

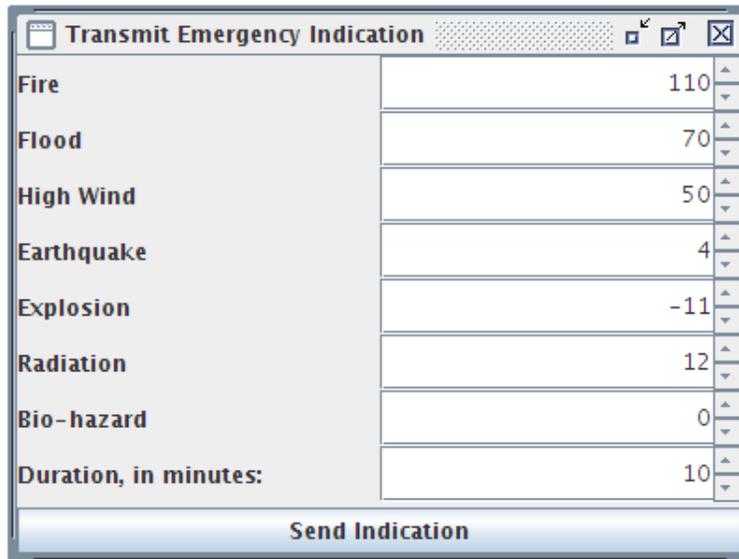


Figure 6.9: Emergency indication dialog.

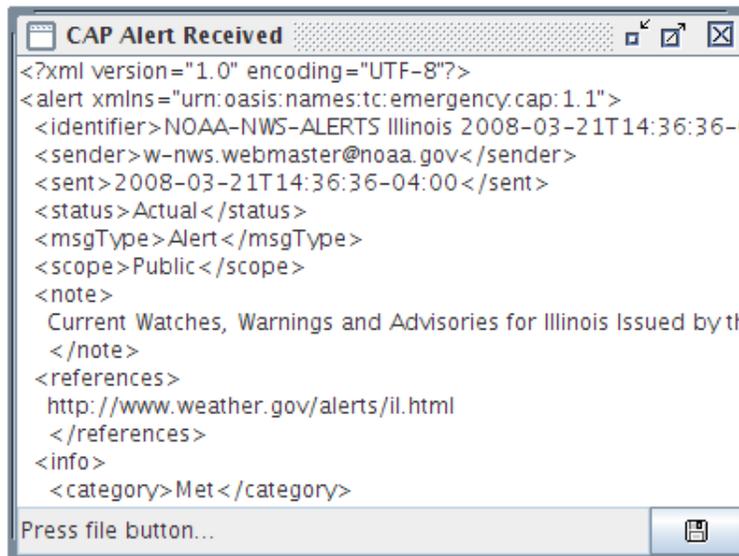


Figure 6.10: CAP alert reception dialog.

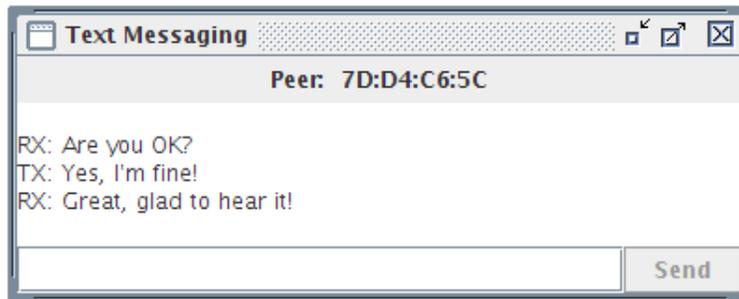


Figure 6.11: Text messaging dialog.

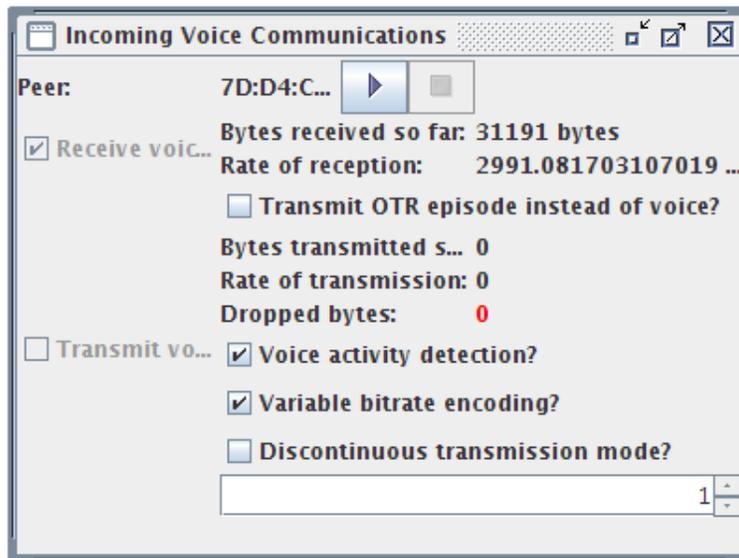


Figure 6.12: Audio messaging dialog.

number of nodes, the CAP message format may be suitable. CAP messages can be loaded and then transmitted, and are then displayed on recipient nodes as shown in Figure 6.10.

Many ERN services are intended to operate between specific pairs of nodes, in a session-oriented fashion. These can be accessed using the “Create Comm Session” button on the testbench. Three types of messaging data are supported: Text, audio, and images. When a text messaging session is established between two nodes, a dialog is shown on each node’s testbench that resembles a traditional instant messaging client (Figure 6.11).

Audio exchanges also use a bidirectional client dialog on each node, but it has specialized controls to adjust the properties of the audio stream (Figure 6.12). The dialog is

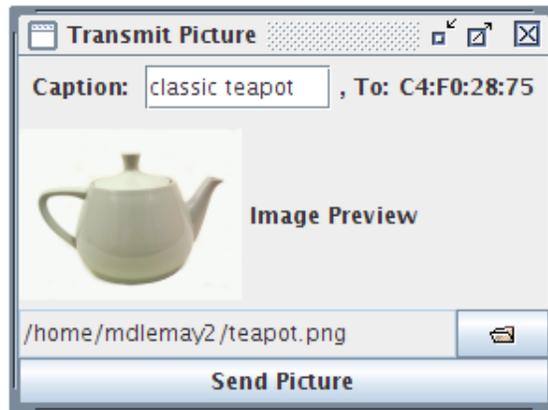


Figure 6.13: Image transmission dialog.

logically divided into two sections. The upper section is dedicated to audio reception, while the lower controls audio transmission. Both sections indicate the total amount of data received/transmitted and the average throughput. The transmission section also indicates how many packets have been dropped. The play and stop buttons at the very top of the dialog control both transmission and reception, but those subsystems can be independently controlled using the checkboxes on the left side of the dialog. The transmission section of the dialog has two main subsections. The first comprises a single control, which determines whether a sample wave file will be transmitted, or alternately whether audio will be recorded in realtime from a microphone or line-in. The second section controls the audio encoding process. The Speex encoder has several parameters that can be adjusted to affect performance. Voice activity detection automatically inserts low-bandwidth filler noise into the audio stream, to reduce bandwidth requirements. Variable bitrate encoding analyzes acoustic properties of the stream and dynamically adjusts the quality of the stream to achieve a better tradeoff between bandwidth consumption and audio fidelity. Discontinuous transmission mode is similar to voice activity detection in that it reduces bandwidth consumption during pauses in speech. If variable bitrate encoding is disabled, the quality of the audio stream can be manually adjusted using a spinner control. Higher quality levels require greater bandwidth.

Image messages are handled asymmetrically in that the sender's dialog box contains controls to select an image file and enter its caption, preview the image prior to sending it, and then send an image message containing the image data (Figure 6.13), while the recipient's dialog box simply displays the image and its caption.

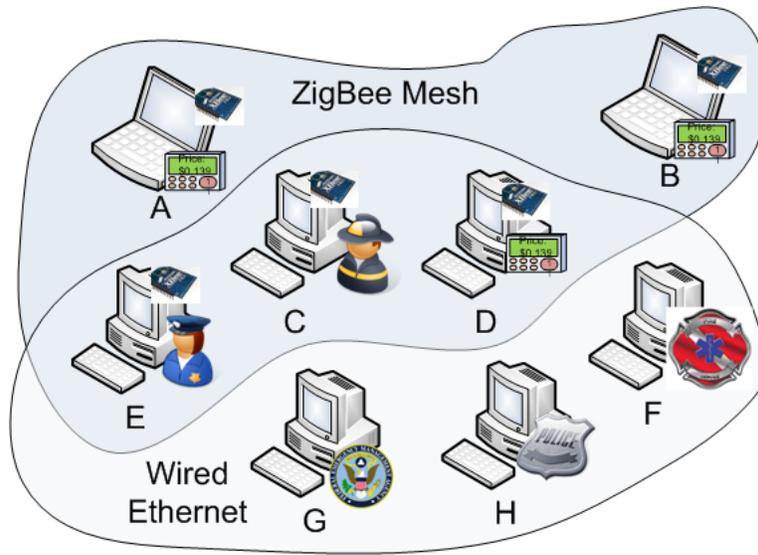


Figure 6.14: Topology of implementation testbed.

6.2 Testbed Characteristics

For the following experiments we constructed a testbed network using five ZigBee nodes each connected to a different host computer. Three of the nodes represent PCTs (A, B, D), and two of them represent rescuers. One of the rescuers is a fire-fighter (C), and the other is a police officer (E). None of these nodes is permitted to route ERN messages (besides emergency indications) on the network until after the network enters emergency-response mode. We also include five Ethernet-connected nodes, two of which are a part of the set of five ZigBee-connected nodes. The remaining three represent a fire chief (F), police chief (H), and FEMA official (G). This network topology is shown in Figure 6.14.

The nodes are all located within a single lab, which produces significant interference between the ZigBee wireless nodes. Thus, our experiments can be expected to provide relatively poor (although not necessarily worst-case) performance results. The two laptop computers ran Windows XP with Sun Java 1.6, while the remaining systems ran Fedora 7 or 8 with Sun Java 1.6. Various processors were in use, but all were Intel Core Duo or Core 2 Duo models, with the exception of node D, which is a Pentium 4.

We used this testbed to evaluate the functionality and verify the proper operation of our system. We also used it to quantify the performance of the AODV layer operating on heterogeneous underlying networks, to determine how useful such a network's

services may be in an emergency. We describe the results of these experiments in the next section.

7 Evaluation

To measure the throughput and latency of the network when heavily-loaded, we used a modified version of the ERN image transfer protocol. We modified the user interface of that function so that it is possible for users to specify how many consecutive times each image should be transmitted across the network. We used an 11148-byte image, which resulted in 11203-byte ERN packets when fully encapsulated.

First, to determine the optimal performance of our prototype, we transmitted 30 packets from node D to node G, over a single Ethernet hop. The throughput and latency of each image packet is depicted in Figure 7.1. The average latency was 620 milliseconds, with a standard deviation of 44.1 milliseconds. The average throughput was 18147.8 bytes per second, with a standard deviation of 1219.7 bytes per second. The latency is calculated for the entire image packet, which is segmented into 8 UDP packets. The latency is computed from the time that the packet is segmented until the time that all of the segments have been reassembled on the receiving end. In our subsequent experiments with ZigBee networks, each image packet is segmented into 163 ZigBee packets. The latency of the UDP/IP packets increases quickly as the first few packets are transmitted, and then returns to a lower level. It may be possible to explain this phenomenon as being caused by subtleties of the underlying operating systems' UDP implementations.

Next, we transferred the images from node A to node G, via node D. Node A only has ZigBee connectivity, and node G only has wired Ethernet connectivity, while node D has both. Thus, each packet traversed a total of 2 hops, the first over ZigBee, and the second over multicast UDP/IP. Thirty total packets were transmitted. The throughput and latency of each image packet is depicted in Figure 7.2. The average latency was 74.7 seconds, with a standard deviation of 60.8 seconds. The average throughput was 397 bytes per second, with a standard deviation of 466.6 bytes per second. To achieve reasonable performance from the XBee modem, it is necessary to support up to 128

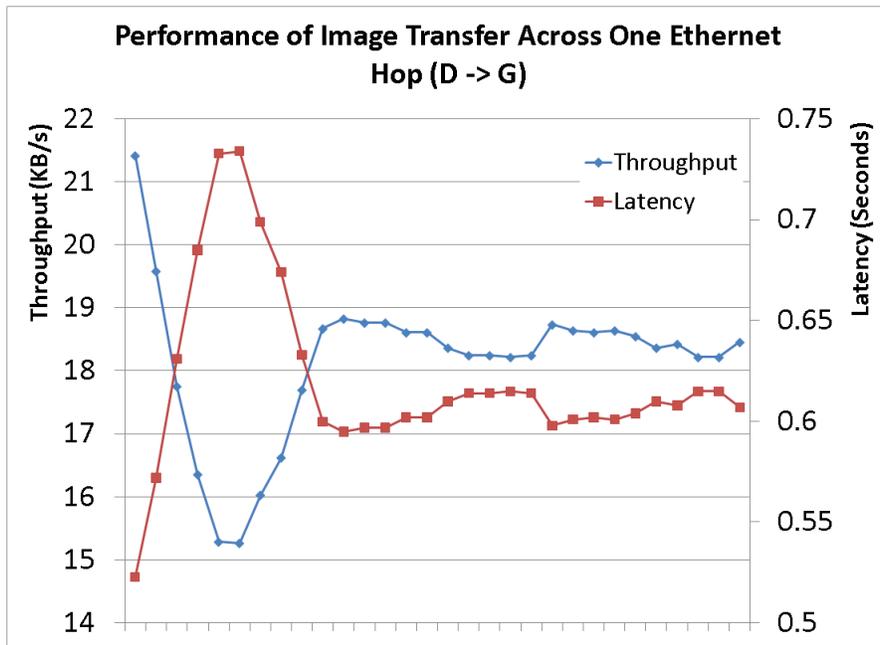


Figure 7.1: Packet latency and throughput observed during image transfer experiment across one UDP/IP hop.

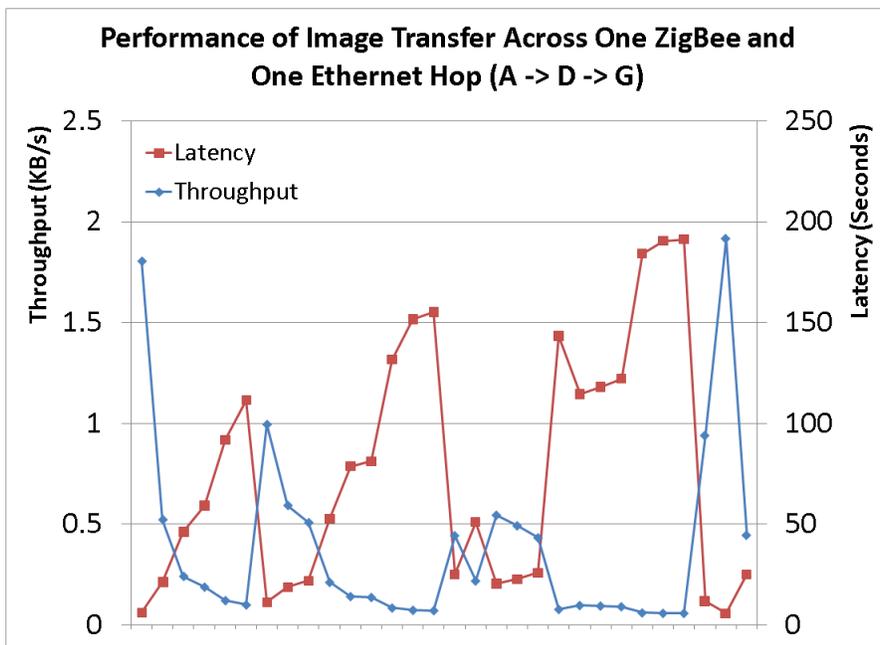


Figure 7.2: Packet latency and throughput observed during image transfer experiment across one ZigBee hop and then one UDP/IP hop.

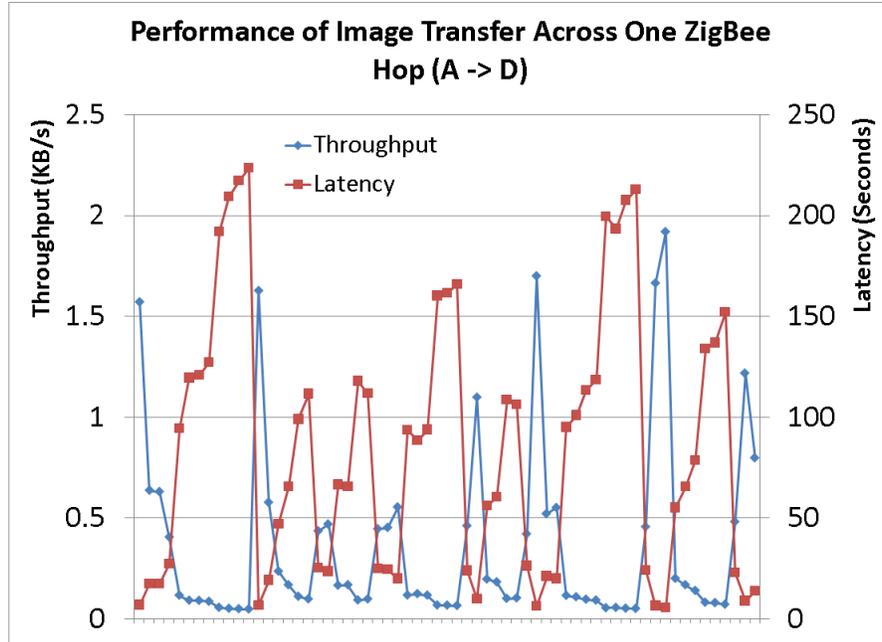


Figure 7.3: Packet latency and throughput observed during image transfer experiment across one ZigBee hop.

outstanding ZigBee packets at any point in time. Thus, many of the latency measurements include some time spent in a queue. The latency measurements exhibit a cyclical pattern that is difficult to explain without a detailed understanding of the implementation of XBee modems. The cycles may be due to a collision-response algorithm, since collisions clearly interrupted the transmissions. The collisions in this case occurred after the reassembly of each packet, at which time the recipient responded with an ack message. Upon receiving an ack message while simultaneously transmitting the next image packet, the sending XBee modem halted its operation for more than a second before resuming transmissions. This could either be due to the ZigBee specification itself, a poor implementation on the part of Maxstream, or an undiscovered bug in our prototype.

The cyclical pattern persists even if a larger number of image packets are transmitted, as is clear from Figure 7.3, which summarizes the performance results of our final successful experiment, transmitting image packets across a single ZigBee hop. We transmitted 63 packets from node A to node D. The average latency was 85.1 seconds, with a standard deviation of 65.9 seconds. The average throughput was 372.5 bytes per second, with a standard deviation of 467.0 bytes per second.

We also attempted to perform simultaneous image transfers between independent pairs of nodes, but the ZigBee transmissions apparently collided and caused both connections to cease communicating for seconds at a time, with successful transmissions of single 69-byte segments occurring only infrequently. Thus, it seems clear that either ZigBee is poorly suited for such applications, Maxstream has poorly implemented ZigBee, or there is an undiscovered error in our prototype or its support software leading to this disappointing outcome. Unsurprisingly, all of these experimental results suggest that the ZigBee mesh network is the major performance bottleneck of our prototype ERN. Regardless, we still believe that it is beneficial to use ZigBee mesh networks as ERN segments, since even very limited communication can be useful during emergencies. For example, text messages require far fewer packet transmissions than images, decreasing the likelihood of collisions, but can still serve as a useful communications channel between victims and rescuers.

8 Conclusion

ERNs must be included in any comprehensive emergency-response plan, and are likely to be increasingly important as the benefits they provide during rescue and recovery operations become more widely recognized. To increase the robustness and coverage of ERNs, we have proposed that robust commercial and governmental networks such as those supporting AMI and surveillance cameras be integrated with other, dedicated ERNs. It is critically important that ERNs be properly provisioned at all times, since lives may be directly dependent on the proper functioning of emergency-response applications. Thus, we adapted a network provisioning algorithm to design and analyze emergency-response networks with respect to various levels of hazard exposure and availability. We recommended modifications to devices implementing ERNs to provide the necessary support for the network, and also specified simple network protocols that can be used to provide inter-operable ERN services on heterogeneous networks. Since the routing protocols of the underlying networks may be incompatible with each other, or unsuitable for ERN scenarios, we showed that the AODV routing protocol used in mobile ad-hoc networks can support ERN services on popular types of networks. We evaluated the performance of a testbed network based on a Java prototype implementation of our protocols, and highlighted some of the performance difficulties that can arise when using a specific brand of ZigBee modems in ERNs. In contrast, the performance of our testbed on a wired Ethernet network was perfectly acceptable.

References

- [ABB⁺04] W. Aiello, S.M. Bellovin, M. Blaze, R. Canetti, J. Ioannidis, A.D. Keromytis, and O. Reingold. Just Fast Keying: Key agreement in a hostile internet. *ACM Transactions on Information and System Security*, 7(2):242–273, 2004.
- [ACVS02] G.S. Ahn, A.T. Campbell, A. Veres, and L.H. Sun. SWAN: Service differentiation in stateless wireless ad hoc networks. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '03)*, volume 2, pages 457–466, April 2002.
- [Ass06] Association of Public Television Stations. Digital emergency alert system (DEAS) fact sheet. http://www.fema.gov/pdf/media/2006/deas_fact_sheet.pdf, July 2006.
- [BJR02] S. Borenstein, M. Jaske, and A. Rosenfeld. Dynamic pricing, advanced metering and demand response in electricity markets. *Center for the Study of Energy Markets*, October 2002.
- [BOR02] J. Bram, J. Orr, and C. Rapaport. Measuring the effects of the September 11 attack on New York City. *Federal Reserve Bank of New York (FRBNY) Economic Policy Review*, November 2002.
- [Bot03] A. Botterell. An advanced EAS relay network using the common alerting protocol (CAP). http://www.incident.com/cap/docs/aps/Advanced_EAS_Concept.pdf, September 2003.
- [CH06] Louise K. Comfort and Thomas W. Haase. Communication, coherence, and collective action: The impact of Hurricane Katrina on communications infrastructure. *Public Works Management Policy*, 10(4):328–343, April 2006.
- [CHKS06] C. Chekuri, M.T. Hajiaghayi, G. Kortsarz, and M.R. Salavatipour. Approximation algorithms for non-uniform buy-at-bulk network design. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '06)*, pages 677–686, Washington, DC, USA, October 2006.
- [CL00] Y.S. Chen and K.C. Lai. MESH: multi-eye spiral-hopping routing protocol in a wireless ad hoc network. In *Proceeding of the 9th International Conference on Computer Communications and Networks (ICCCN '00)*, pages 657–661, October 2000.
- [Com] Federal Communications Commission. 700 MHz public safety spectrum. <http://www.fcc.gov/pshs/public-safety-spectrum/700-MHz/>.
- [Con04] EPRI IntelliGrid Consortium. Automatic meter reading (AMR) and related customer service functions. *IntelliGrid Architecture Use Cases*, 2004.

- [DKG⁺04] L. Ding, P. Kolari, S. Ganjugunte, T. Finin, and A. Joshi. Modeling and evaluating trust network inference. In *Proceedings of the 7th International Workshop on Trust in Agent Societies at AAMAS '04*, July 2004.
- [Ebe04] D. Eberhart. FCC's emergency alert system coming to your cell phone soon. <http://archive.newsmax.com/archives/articles/2004/8/24/210743.shtml>, August 2004.
- [Ega05] D. Egan. The emergence of ZigBee in building automation and industrial control. *Computing & Control Engineering Journal*, 16(2):14–19, 2005.
- [Els05] Elster Electricity, LLC. Salt river project moves forward with production-scale rollout of elster electricity's EnergyAxis system. http://tdworld.com/distribution_management_systems/power_salt_river_project_3/, February 2005.
- [Fau07] G.R. Faulhaber. Solving the interoperability problem: Are we on the same channel? An essay on the problems and prospects for public safety radio. *Federal Communications Law Journal (Indiana University School of Law-Bloomington)*, 59(3):493, 2007.
- [FDW06] C.M. Firestone, E. Director, and P.J. Weiser. Clearing the air: Convergence and the safety enterprise. *Communications and Society Program (The Aspen Institute)*, May 2006.
- [FRL05] C.L. Fok, G.C. Roman, and C. Lu. Rapid development and flexible deployment of adaptive wireless sensor network applications. In *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems (ICDCS '05)*, pages 653–662, June 2005.
- [Gov06] Government Technology. Beijing deploys wireless mesh network in preparation for 2008 Olympics. http://www.govtech.net/magazine/channel_story.php/100646, August 2006.
- [Gre06] T. Greene. New Orleans' Wi-Fi network now a lifeline. <http://www.computerworld.com/mobiletopics/mobile/story/0,10801,109662,00.html> (*Computerworld Mobile/Wireless*), March 2006.
- [HABR05] K.A. Harras, K.C. Almeroth, and E.M. Belding-Royer. Delay tolerant mobile networks (DTMNS): Controlled flooding in sparse mobile networks. *IFIP Networking*, 3462/2005:1180–1192, 2005.
- [Har05] B. Harris. ZigBee joins Wi-Fi as AMR alternative. <http://www.govtech.com/gt/articles/97681> (*Government Technology*), December 2005.
- [HBZ⁺06] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: A distributed mobile sensor computing system. In *Proceedings of the 4th international Conference on Embedded Networked Sensor Systems (SenSys '06)*, pages 125–138. ACM Press New York, NY, USA, October 2006.
- [HKH05] D. Hinton, T.E. Klein, and M. Haner. Emergency response networks with broadband services. *Bell Labs Technical Journal*, 10(2):121–138, 2005.

- [HMA⁺99] Michael Hicks, Jonathan T. Moore, D. Scott Alexander, Carl A. Gunter, and Scott Nettles. PLANet: An active internetwork. In *Proceedings of the 18th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '99)*, pages 1124–1133, March 1999.
- [Itr06] Itron, Inc. Manitoba hydro first to deploy new advanced metering and communication technology from Itron. http://www.itron.com/pages/news_press_individual.asp?id=itr_008600.xml, June 2006.
- [Kin03] P. Kinney. ZigBee technology: Wireless control that simply works. <http://hometoys.com/htinews/oct03/articles/kinney/zigbee.htm> (*Home Toys*), October 2003.
- [KK00] B. Karp and H. T. Kung. Gpsr: greedy perimeter stateless routing for wireless networks. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MobiCom '00)*, pages 243–254, August 2000.
- [LAN03] LAN/MAN Standards Committee. IEEE standard for information technology – 802.15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (LR-WPANs). *IEEE Computer Society*, May 2003.
- [LMFJ⁺04] K. Lorincz, DJ Malan, TRF Fulford-Jones, A. Nawoj, A. Clavel, V. Shnayder, G. Mainland, M. Welsh, and S. Moulton. Sensor networks for emergency response: Challenges and opportunities. *IEEE Pervasive Computing*, 3(4):16–23, 2004.
- [mat06] mattw (psuedonymous member at instructables.com). Mobile emergency communications: Mobile repeater and mesh node. <http://www.instructables.com/id/ERXV7WVTB8EP2872PP/>, April 2006.
- [MB02] S.F. Midkiff and C.W. Bostian. Rapidly-deployable broadband wireless networks for disaster and emergency response. *Presented at The First IEEE Workshop on Disaster Recovery Networks (DIREN 02)*, June 2002.
- [Moo05] L.K. Moore. Public safety communications: Policy, proposals, legislation and progress. *DTIC Research Report ADA453736*, June 2005.
- [Moo06] L.K. Moore. Emergency communications: The emergency alert system (EAS) and all-hazard warnings. *CRS Report for Congress*, July 2006.
- [Pea88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [PR99] C.E. Perkins and E.M. Royer. Ad-hoc on-demand distance vector routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '99)*, volume 2, pages 90–100, February 1999.
- [PS00] J.S. Park and R. Sandhu. Binding identities and attributes using digitally signed certificates. *Proceedings of the 16th Annual Computer Security Applications Conference (ACSAC '00)*, page 120, December 2000.
- [Res05] Chartwell Market Research. Mesh technologies likely to drive utilities to deploy wireless AMR technologies, says new chartwell report. http://www.chartwellinc.com/pressrelease.cfm?pressrelease_id=29, January 2005.

- [sce07] Southern california edison achieves key advanced metering goal. <http://electricenergyonline.com/IndustryNews.asp?m=1&id=71649> (*Electric Energy Online*), August 2007.
- [SD05] R. Sharma and M. Dillon. ZigBee-based utility meters help china handle massive housing boom. http://www.zigbee.org/imwp/idms/popups/pop_download.asp?contentID=7317, December 2005.
- [Str] Strix Systems. Strix Systems press releases. <http://www.strixsystems.com/events/default.asp>.
- [TG97] K.J. Tierney and J.D. Goltz. Emergency Response: Lessons Learned from the Kobe Earthquake. <http://www.udel.edu/DRC/preliminary/260.pdf>, 1997.
- [Vai02] N.H. Vaidya. Weak duplicate address detection in mobile ad hoc networks. In *Proceedings of the 3rd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc '02)*, pages 206–216, June 2002.
- [Val06] J.M. Valin. Speex: A Free Codec For Free Speech. http://people.xiph.org/~jm/papers/speex_lca2006.pdf, 2006.
- [Was08] CNN Washington. Cell phone alert system proposed. http://money.cnn.com/2008/04/09/technology/fcc-cell_phone_alert, April 2008.
- [Zig05] ZigBee Alliance. ZigBee security specification overview. <http://www.zigbee.org/en/events/documents/December2005-Open-House-Presentations/ZigBee-Security-Layer-Technical-Overview.pdf>, December 2005.
- [Zig06] ZigBee Alliance. ZigBee specification. <http://www.zigbee.org>, 2006.

Author's Biography

Michael David LeMay received his B.S. in Computer Science from the University of Wisconsin-Eau Claire in 2005. He has interned with the U.S. Department of Defense and Cray, Inc. He is expecting his M.S. in May 2008 from the University of Illinois at Urbana-Champaign (UIUC). Currently, he is a PhD student in Computer Science at the University of Illinois at Urbana-Champaign and working as a research assistant with Professor Carl A. Gunter in Illinois Security Lab. He has been awarded the NDSEG Fellowship supporting him from August 2005 until August 2008. He has publications in the ACM Symposium on Access Control Models and Technologies, ACM Workshop on Hot Topics in Networking, Hawai'i International Conference on System Sciences, and several other forums.