Supporting Emergency-Response by Retasking Network Infrastructures*

Michael LeMay and Carl A. Gunter
University of Illinois at Urbana-Champaign

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 [2]. Even more significant disruptions to communication channels occurred when Hurricane Katrina rendered most of the infrastructure components within its wake partially or completely inoperable [4]. 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 [14]. 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 every year.

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 [7]. 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 [11]. Buildings are also being enhanced with mesh networks for building automation [5].

Mesh networks are more resilient to node failures than other types of networks, which makes them a logical choice for such applications. However, like any other infrastructure improvement, it is often expensive to deploy mesh 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 mesh network carries only government communications, the level of value may be relatively low compared to commercial networks, ultimately causing the governmental mesh networks to be smaller than commercial networks. 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 [10]. 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 path is the retasking of existing networks for ERN 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.

How is it possible to use networks for emergency communications if they were originally intended to provide a different service while also ensuring that the original service is not adversely affected during ordinary operations? There are at least three primary considerations in answering this question: *detection*, *platform support*, and *topology*. 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. 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

*Published in the Proceedings of the 6th ACM Workshop on Hot Topics in Networking (HotNets-VI), November 2007.

their locations may be unpredictable. The aim of this paper is to discuss these three challenges with respect to technologies appropriate to ERN retasking. Our primary technical proposal is using buy-at-bulk network provisioning algorithms as a strategy for assessing the readiness and cost effectiveness of network infrastructure for ERN retasking.

We use Advanced Metering Infrastructure (AMI), surveillance camera mesh networks, and enhanced cellphones as running examples throughout this paper. AMI networks are a relatively recent development, so we provide some background on them here. Traditional mechanical electric meters are rapidly being replaced by microcontroller-based electronic meters that are often connected to each other and the central MDMA using robust mesh networks. They provide many business benefits, such as the ability to use more advanced power pricing scheme such as Real-Time Pricing (RTP), instant outage detection, and demand response [1]. These benefits are driving wide deployments of advanced meters. Thus, if these networks could be adapted to support ERN, they could greatly increase coverage with relatively small additional investments.

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.

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 for long distances without additional infrastructure. Techniques for establishing emergency-response backbones are presented in [13], and delay-tolerant networking may also be useful in certain situations [8]. Enhanced cellphones could be particularly helpful in this regard, because their radios are relatively powerful and could greatly extend the coverage of an ERN. The three primary device types we are considering also offer vastly different user interfaces. AMI networks are often required to have linkages to building- and home-automation networks, most commonly to smart thermostats using ZigBee (described later). These smart thermostats provide a convenient user interface near each node, which can be useful in ERN scenarios. 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. Of course, cellphones include very sophisticated integrated user interfaces. 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 Challenges and Requirements

Detection Determining when emergency conditions are in place may be important from an access control standpoint. Some devices that could provide emergency-response networking are equipped with sensors that may help to determine when an emergency has occurred. For example, if an advanced meter detects a power outage, that may serve as an indication that a significant emergency may have occurred. However, such indications are typically ambiguous and may not be sufficient to support a declaration that an emergency is actually in effect. Thus, it may be necessary to also rely upon local and remote external inputs, such as human inputs and indications from the central network operator. In the case of an advanced meter, it may have a connection to a smart thermostat equipped with a "panic button" that permits a human occupant to signal that an emergency has occurred and assistance from a rescuer is requested. Surveillance cameras are not equipped with local interfaces and are often completely inaccessible to humans, but they could potentially be enhanced with emergency-detection sensors that respond to signals for emergency personnel and also listen for indications from the network operator when the connection is available.

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. Unfortunately, humans may have other motives for signaling an emergency, particularly if they wish to use the ERN for illegitimate purposes or during non-emergency conditions, and such manipulation may be possible when the connection to the network operator is severed. We must prevent such compromises from occurring, since that would adversely affect the primary functions of the network. Threshold schemes may be useful in this case, so that each device only provides ERN connectivity if some threshold of emergency indications from human users in the device's vicinity have been received, where that threshold may be dynamically adjusted based upon sensor readings and other factors. For example, the network operator may monitor weather conditions, terrorist threat alert levels, and other emergency condition predictors, and broadcast future threshold adjustments to victim devices before the adjustments need to occur. More complicated mechanisms could be devised, but could potentially introduce unacceptable delays between the onset of emergency conditions and the availability of ERN services.

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. 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 incapable 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 WiFi 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.

Once ERN services have been activated, the ERN must support communications 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.

Platform Support To maximize the utility of ERNs, standard emergency-response applications must be developed. If an ERN is being used to support existing applications such as those that use IP, this is not an issue. However, if specialized requirements for the network exist, such as the requirement that it be used to locate survivors within a disaster zone as efficiently as possible, existing applications are unlikely to suffice. Two distinct responses are plausible in this situation. 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 [10].

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* [9] aim to add software to network elements and there have been efforts to provide dynamically extensible platforms for sensors [6].

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.

ZigBee is a wireless protocol stack for low-rate wireless personal-area networks [15]. 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 [12], 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.

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 preinstalled 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.

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 attacks can produce electro-magnetic 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.

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 are 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 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.

Topology 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. Emergency-response network 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.

First, we represent ERNs as graphs. Essentially, emergency-response networks comprise 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.

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 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 1. 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 $\mathcal{T} = (s_1, t_1), \ldots, (s_h, t_h)$, where $s_i, t_i \in V$ and each pair is associated with a non-negative bandwidth demand d_i.

3 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 topology 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.

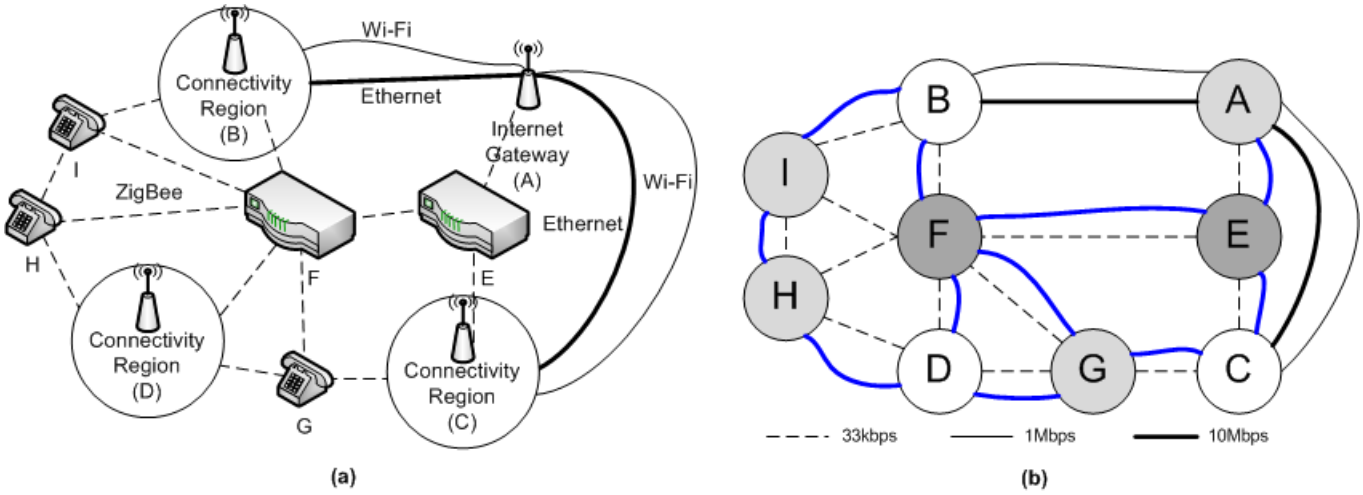

Figure 1: *a)* Initial network provisioning graph, *b)* Simplified network provisioning graph with sources (light gray), sinks (white), and routers (dark gray) identified. Also includes provisional links in blue.

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 E' edges 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 [3] 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 [3]. 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 impor-

tant 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 unfeasibly high 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.

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 $c_e = (\textit{equipcost}(e) + \textit{lw} \cdot \textit{latency}(e))/\textit{reliability}(e)$, where $\textit{equipcost}(e)$ is the cost to install and maintain e, $\textit{lw}$ 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 $\textit{reliability}(e) = \sum_{d \in \mathcal{D}}(p(e,d) \cdot \textit{dependability}(e,d))$, 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:

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

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, $\textit{susceptibility}(e,h)$ is the probability that h will degrade or destroy edge e, and $\textit{availability}(e,h)$ 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:

$$\begin{aligned} l_e \;=\; & \max\Big(l_{\textit{nextbigger}(e)}, \Big(\textit{bwchrg}(e) \\ & + \max\Big(0, \Big(\frac{(c_{\textit{nextbigger}(e)} - c_e)}{\textit{capacity}(e)} - \textit{bwchrg}(e)\Big)\Big)\Big)\Big), \end{aligned}$$

where $\textit{nextbigger}(e)$ is the edge that is parallel to e and has the next lowest bandwidth, $\textit{capacity}(e)$ is the bandwidth of e, and $\textit{bwchrg}(e)$ 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 [3]. 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.

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 preinstalled 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 routing algorithms more akin to those used in ad-hoc networks and ZigBee may be more appropriate. The mechanism for supporting such algorithms on heterogeneous IP networks is an important research problem.

4 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 rebuilding operations become more widely recognized. To increase the robustness and coverage of ERNs, we have proposed that robust networks such as those for AMI 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.

Acknowledgements

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.

References

[1] S. Borenstein, M. Jaske, and A. Rosenfeld. Dynamic Pricing, Advanced Metering and Demand Response in Electricity Markets. *The Energy Foundation*, pages 10–11, 2002.

[2] 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.

[3] C. Chekuri, M. T. Hajiaghayi, G. Kortsarz, and M. R. Salavatipour. Approximation algorithms for non-uniform buy-at-bulk network design. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 677–686, Washington, DC, USA, 2006. IEEE Computer Society.

[4] L. K. Comfort and T. W. Haase. Communication, coherence, and collective action: The impact of Hurricane Katrina on communications infrastructure. *Public Works Management and Policy*, 10, April 2006.

[5] D. Egan. The emergence of ZigBee in building automation and industrial control. *Computing & Control Engineering Journal*, 16(2):14–19, 2005.

[6] C. Fok, G. Roman, and C. Lu. Rapid Development and Flexible Deployment of Adaptive Wireless Sensor Network Applications. *Distributed Computing Systems, 2005. ICDCS 2005. Proceedings. 25th IEEE International Conference on*, pages 653–662, 2005.

[7] T. Greene. New Orleans' Wi-Fi network now a lifeline. *Computerworld Mobile/Wireless*, March 2006.

[8] K. Harras, K. Almeroth, and E. Belding-Royer. Delay tolerant mobile networks (dtmns): Controlled flooding in sparse mobile networks. *IFIP Networking*, 2005.

[9] M. Hicks, J. T. Moore, D. S. Alexander, C. A. Gunter, and S. Nettles. PLANet: An active internetwork. In *Conference on Computer Communications (INFOCOM '99)*, pages 1124–1133, Boston, MA, March 1999. IEEE.

[10] D. Hinton, T. Klein, and M. Haner. Emergency Response Networks with Broadband Services. *Bell Labs Technical Journal*, 10(2):121–138, 2005.

[11] G. Johnston. Mesh technologies likely to drive utilities to deploy wireless AMR technologies, says new chartwell report. http://www.primenewswire.com/newsroom/news.html?d=71958, Jan. 2005.

[12] 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.

[13] S. Midkiff and C. Bostian. Rapidly-Deployable Broadband Wireless Networks for Disaster and Emergency Response. *Presented at The First IEEE Workshop on Disaster Recovery Networks (DIREN 02)*, 2002.

[14] K. Tierney and J. Goltz. Emergency Response: Lessons Learned from the Kobe Earthquake. http://www.udel.edu/DRC/preliminary/260.pdf, 1997.

[15] ZigBee Alliance. ZigBee specification. http://www.zigbee.org, 2006.