WIISARD: A Measurement Study of Network Properties and Protocol Reliability during an Emergency Response

Octav Chipara[1], William G. Griswold[2], Anders N. Plymoth[3], Ricky Huang[3], Fang Liu[3],
Per Johansson[3], Ramesh Rao[3], Theodore C. Chan[4], Colleen Buono[4]
[1]Department of Computer Science, University of Iowa
[2] Department of Computer Science, University of California San Diego
[3] CalIT2, University of California San Diego
[4] Division of Emergency Medicine, Department of Medicine, University of California San Diego

ABSTRACT

This paper describes the design, deployment, and empirical evaluation of WIISARD – a novel emergency response system that provides reliable communication in dynamic wireless environments without extensive communication infrastructure. The main contribution of this paper is an in-depth empirical study of network properties that emerge during a drill in which WIISARD is deployed with minimal infrastructure support. The drill involves 19 first responders and 41 victims. The properties of links established among first responders vary between phases of the drill and depend upon the responder's role in the drill. The rescue phase – in which responders are highly mobile as they triage victims – poses significant challenges to reliable communication. During this phase, the contacts between responders are short-lived; however, they are reestablished within minutes. Once a contact between responders is established, the quality of the link between those responders is usually high. The connectivity graph observed during the rescue phase is usually connected and has a small diameter although there are times when it has a large diameter or it is partitioned. While mobility increases network dynamics, we also observe that the mobility patterns characteristic of the emergency response workflow can be leveraged to disseminate data efficiently through data muling. WIISARD employs a gossip-based protocol and supports data dissemination through local communication and data muling to achieve 98% reliability during the drill exercise. These results indicate the feasibility of providing reliable communication in emergency response with minimal infrastructure in spite of network dynamics.

Categories and Subject Descriptors

C.2.4 [**Computer-communication Networks**]: Distributed Systems

MobiSys'12, June 25–29, 2012, Low Wood Bay, Lake District, UK.
Copyright 2012 ACM 978-1-4503-1301-8/12/06 ...$10.00.

Keywords

Emergency response, Delay tolerant networking, Mobility, Reliability

1. INTRODUCTION

Every year, natural and human-caused disasters in the form of hurricanes, earthquakes, and infrastructure failures put the health of millions at risk. The key to successful emergency responses is to achieve *situational awareness* through effective communication: information about the event, casualties, and available resources must be exchanged among first responders in a *reliable* and *timely* manner. A breakdown in communication can result in ineffective use of the limited resources or jeopardize the safety of responders.

A typical emergency response involves four overlapping phases: staging, rescue, treatment, and transport. During the staging phase responders arrive on scene and establish command and control. In the rescue phase, the responders triage victims and provide minimal medical care to stabilize their condition. As the scale of the disaster increases, opportunities arise during the treatment phase to provide further, more comprehensive, medical treatment. In the transport phase, the victims are transported to hospitals based on the severity of their injuries.

Traditionally, information during emergency responses has been collected and exchanged either in writing or verbally using hand radios. While written and verbal communication is adequate for incidents that involve a few victims (e.g., car accidents), their effectiveness deteriorates rapidly as the incidents increase in scale. In large-scale incidents, information captured on paper can be lost, while verbal communication often introduces inaccuracies [24]. Moreover, neither form of communication is effective in sharing information rapidly among numerous first responders.

The aim of the WIISARD project is to develop a reliable communication infrastructure for emergency response by taking advantage of mobile computing technology. Capturing information in digital form and exchanging it wirelessly ought to reduce inaccuracies, limit accidental data loss, and facilitate the timely and effortless dissemination of information among first responders. Yet, to meet this goal, we must overcome three key challenges: (1) minimize the reliance on network infrastructure during emergency responses, (2) cope with a dynamic radio environment subject to interference, and (3) support communication among mobile users.

Existing emergency response systems may be divided into two classes based on the networking technology they use. Cellular networks are used increasingly to deliver patient information during their transport to hospitals [29,33]. However, cellular networks are not a reliable solution for emergency response systems as indicated by recent disasters (e.g., Katrina) where the cellular infrastructure either failed or was overloaded. A promising alternative is to adopt mesh networking technology: first responders would bring a number of wireless nodes to the scene that, when deployed, would self-organize in a mesh network that facilitates communication at the disaster scene. Typically, these systems would use either ad-hoc routing (e.g., OLSR, AODV) or Delay Tolerant Networking (DTN) protocols to share information.

The initial design of WIISARD [6] employed a client-server architecture and AODV-based routing [38]. Unfortunately, this design performed inconsistently across deployments. This experience is consistent with the results reported by other researchers [12] and highlights the acute need to develop an empirical characterization of the network properties observed during emergency responses in order to provide a sound basis for understanding the challenges of reliable communication in such a setting. While numerous emergency response systems have been designed and deployed, these systems seldom provide a detailed study of network properties and reliability.

This paper makes the following contributions: (1) We present one of the first empirical studies of network properties that emerge during a drill exercise. The novel aspect of our drill is that it evaluates the effectiveness of using DTN to support reliable communication during emergency responses when minimal infrastructure is deployed. Based on these results, we provide insights regarding the challenges of reliable communication in emergency response and discuss approaches to cope with these challenges. (2) WIISARD achieved 98.25% reliability in spite of significant network dynamics due to responder mobility and network partitions by using a peer-to-peer architecture in conjunction with a gossip-based communication protocol. (3) In a wider context, our work contributes to the growing body of empirical studies of the properties of DTN applications and human mobility patterns. In contrast to these previous studies, a unique aspect of emergency response systems is that responders *cooperate* to rescue victims according to a pre-established workflow (see Section 3). Therefore, we observe more complex network properties that vary with the role of the responders and among the phases of the drill.

We deployed WIISARD as part of a drill organized at University of California, San Diego (UCSD) on August 10th, 2011. The drill involved 19 first responders who rescued 41 victims. The collected data indicates that each stage of the drill has different network properties. The staging and treatment phases exhibited good and stable connectivity. In contrast, the rescue phase posed significant challenges to reliable communication. During this phase, the contacts between providers, which are indicative of the long-term link quality, tended to be short-lived; however, these contacts were reestablished within minutes. Once a contact between responders was established, the associated link was usually characterized by good short-term link quality, as indicated by a median packet reception rate that exceeds 70%.

The connectivity graph of responders during the rescue phase was typically dense and had a small diameter (less than 3 hops). However, due to limited infrastructure, at times the graph became sparse and had a large diameter (7 hops). Even worse, for 26% of the rescue phase, the connectivity graph was partitioned. The mobility of responders was a major source of variability in network properties. However, mobility also had a beneficial effect in that it improved data dissemination through data muling. In fact, due to the inherent mobility patterns of the emergency response workflow, a responder encountered all other responders within seven minutes. This indicates the potential of improving reliability through data muling.

The redesigned WIISARD system that we describe in this paper has a peer-to-peer architecture and data is disseminated through a gossip-based protocol called **W**IISARD **C**ommunication **P**rotocol (WCP). WCP uses local communication to handle variations in link quality and caches data aggressively to support data dissemination via data muling. WCP achieved 98% reliability in spite of significant variations in network properties discussed above. Thus, it is feasible to achieving high reliability in emergency responses even with minimal infrastructure.

The remainder of the paper is organized as follows. The next section discusses related work. Background information regarding emergency responses is provided in Section 3. WIISARD's architecture and software components are described in Section 4. The drill exercise is presented in Section 5. We analyze the network properties that emerge during the drill in Section 6. WIISARD's reliability is characterized in Section 7. We discuss the implications of the observed network properties on emergency response systems in Section 8. Conclusions are provided in Section 9.

2. RELATED WORK

In this section, we review the prior work on emergency response systems and place our work in the wider context of DTN empirical studies.

2.1 Emergency Response Systems

Numerous emergency response systems take advantage of mobile computing technology to improve communication accuracy and timeliness. These systems employ various wireless technologies: Wi-Fi [1,10,20,24,39], 802.15.4 [12,20,26], and cellular networks [29, 33]. We focus on systems that use Wi-Fi mesh networks, as they are closely related to our work, and discuss how they address the challenges of reliable communication in emergency responses.

Early emergency response systems adopted client-server architectures due to their simplicity [4, 17, 18, 24, 39]. An important limitation of this architecture is that clients that are within communication range cannot communicate unless they maintain connections to the server, potentially over multiple hops. To remove this limitation, several systems opted for more flexible network architectures that support either multi-cast [10] or publish-subscribe [20] primitives. However, similar to the client-server systems, these systems still require the construction and maintenance of multi-hop end-to-end routes. Emergency responses are highly dynamic as the wireless channel fluctuates due to the movement of people, vehicles, and equipment (see Section 5). Moreover, since infrastructure is often limited, network partitions are common. These factors make it difficult (if not impossible) to maintain end-to-end routes and, as a result, such systems often suffer from poor reliability in realistic deployments.

To address user mobility and network partitions, recent systems [3, 7, 19, 21] adopt decentralized peer-to-peer architectures and employ delay tolerant networking (DTN) techniques. For example, DistressNet [21] has a hierarchical network architecture that uses 802.15.4 to monitor the vital signs of victims/responders and 802.11 for long-range communication between responders. DistressNet copes with network partitions through DTN techniques. Dong et al. present the Emergency Delay Tolerant Network architecture and investigate the use of Session Initiation Protocol (SIP) based communication [19]. WIISARD also takes advantage of DTN techniques to improve data dissemination through data muling. However, in contrast to the above techniques [3, 19, 21], which have been evaluated only through simulations, WIISARD is evaluated through an in-situ deployment.

In spite of the numerous emergency response systems that have been developed, they are seldom evaluated with real users and in real-world deployments. However, there are a few notable exceptions. As part of the AID-N project [20], the functional requirements of emergency response systems were elicited through surveys of first responders. The usability of the AID-N system was evaluated on a 5-point Likert scale. Surveys indicate that the first responders found AID-N to be more effective in tracking victims than standard paper triage tags. The US Army (BMIST system [30]) and Navy (TACMED-CS [37]) have developed mobile systems to provide access to electronic patient records. These systems are designed to store changes to patient records locally and synchronize them when network connectivity is available. Both systems are currently in use. In our own prior work, we studied the impact of (an earlier version of) WIISARD on the emergency response workflow through a drill exercise [24]. WIISARD significantly reduced the rate of missing/duplicate patient identifiers. None of the above empirical studies provide an analysis of network properties and reliability during emergency responses.

In this paper, we present one of the first empirical studies of network properties during emergency responses. The presented results are complementary with those obtained as part of the Code Blue project. CodeBlue [26] was an earlier system aimed at monitoring patient vital signs during a disaster drill using wireless sensors. Unlike WIISARD, CodeBlue employed an ad hoc multicast routing protocol based on ADMR [13], requiring sensor nodes to maintain routing tables and link quality state throughout the network. The CodeBlue disaster drill study [12] demonstrated the challenges of using such complex, stateful protocols in highly dynamic environments with high mobility and poor link quality In contrast, the focus of this paper is to assess the feasibility of using DTN protocols to support reliable communication in emergency responses when minimal infrastructure is deployed.

2.2 Delay Tolerant Networking

DTN techniques have been proposed for delivering packets when contemporaneous end-to-end paths do not exist. DTN protocols may be classified based on their assumptions of user mobility. The most general DTN protocols (e.g., [8, 25]) do not make any assumptions regarding the mobility of users. These protocols aim at maximizing the likelihood of packet delivery and limiting the number of duplicate packet transmissions by deciding which packets are forwarded during node encounters or dropped when a node's storage capacity is reached. Since little is known about the mobility of responders, the WCP protocol used by WIISARD falls into this category. WCP has similarities to Trickle [25]: in both protocols nodes exchange metadata to determine the state that must be updated and suppress packet retransmissions based on overhearing. The novelty of this paper is not WCP but rather the characterization of network properties and mobility patterns that emerge during emergency responses.

DTN protocols that take advantage of the mobility patterns of users have been proposed (e.g., [15, 27]). Fundamental to these protocols is an accurate characterization of human mobility. Initial models of human mobility focused on random walk models, which did not accurately capture many of the properties of human mobility (for a review refer to [9]). Recently, more realistic models of human mobility have been developed based on empirical mobility traces collected using GPS [35] or based on Bluetooth contacts [11]. The macro scale human mobility has been studied using association traces of mobile laptops/PDAs and 802.11 access points (e.g., traces collected at UCSD [28] and Dartmouth [22]) and based on student class schedules [36]. These traces indicate that some popular locations are significantly more likely to be visited leading to networks in which user densities vary widely. Moreover, it has been observed that the distribution of inter-contact times of users is heavy tailed, a characteristic that cannot be reproduced by random walk models [11, 35]. However, a common modeling assumption is that the mobility patterns of users are identical and network properties are time-invariant [8, 11, 34, 35]. Exceptions do exist. For example, mobility models have been enhanced to include social relationships of users [15, 31]. These models have been shown to support efficient message delivery in DTNs [15]. Similarly, in recent study, Hsu et al. considers the time variant nature of mobility patterns.

A characteristic of the traces considered in these studies is that they studied populations of users that do not cooperate to complete tasks. In contrast, the results presented in this paper consider the cooperative task of rescuing victims during an emergency response according to the workflow discussed in Section 3. This leads to a richer set of mobility patterns and network properties. Specifically, we will show that network properties vary with both the role of the responders and the phases of the drill. This differentiates our results from previous studies where user behavior is considered to be uniform and network properties stationary. Moreover, our work also differs from the efforts of integrating social relationships into motion models since, in emergency response, the organizational structure is dictated by the response workflow rather than socialization behavior. Due to these properties, most of the existing models of mobility are not applicable. In fact, the analysis presented in this paper will be guided by an understanding of the emergency response workflow.

3. BACKGROUND

Most US emergency agencies follow the Incident Command System (ICS) protocol [2]. Next, we summarize the relevant concepts and terminology of the ICS.

The geographic layout of a drill is divided into *zones* of different risk to victims and responders (see Figure 2). The *hot zone* includes the incident location and, thus, it is considered to be unsafe. The time victims and responders spend in the hot zone should be minimized. The *warm* zone has a medium risk while the *cold zone* is considered to be safe.

A typical *emergency response workflow* involves the following teams and associated responsibilities. The entry team locates, triages, and evacuates the victims from the hot to the cold zone. In the cold zone, the medical team retriages victims, performs more detailed medical exams, and provides medical care. The transport team manages the arrival of ambulances and the loading of victims to hospitals. The med-com team tracks available hospital resources and allocates victims to hospital beds based on the victim's injuries and resource availability. Patients flow from one team to the other until they are transported to hospitals. For the workflow to scale successfully to mass casualty incidents, the teams are intended to work independently. However, for the system to work efficiently, communication plays a crucial role in ensuring effective hand-offs as victims are transferred between teams and enhancing global communication awareness as captured by victim counts, victim statuses, and resource availability.

The typical implementation of ICS results in functions such as the command, triage, treatment, and transport having defined locations for the duration of the response. Also, access into and out of the hot zone is typically managed with one or more entry points for safety reasons. Consequently, as responders move victims to the cold zone, they are guided through the hot zone checkpoint and all arrive in the triage area. In this way, the independent actions of the teams result in an orderly, coordinated response. As described in the following sections, the mobility patterns characteristic of the emergency response workflow have a profound impact on network properties.

4. WIISARD SYSTEM

This section presents the system architecture, hardware, and software components of WIISARD. Here we focus on the key design decisions we made to support reliable communication in emergency responses.

4.1 System Architecture

WIISARD supports the emergency response workflow by tracking victims during the response and by providing role-tailored interfaces to manage the electronic medical records of patients, the assignment of patients to ambulances, and the availability of hospital resources. WIISARD consists of *triage devices*, *mid-tier devices*, and *command centers.* The entry and medical teams use the *triage devices* to triage patients and manage their electronic medical records. The workflow of the transport and med-com teams is supported via *mid-tier devices.* The *mid-tier devices* implement capabilities for managing ambulances and hospital resources in addition to the functionality already provided by *triage devices.* The incident commander uses a *command center* to get an overview of the progress of the response by monitoring the locations of responders and accessing statistics including victim counts, victim statuses, and resource availability.

WIISARD supports a simple broadcast model: the information generated by a responder is disseminated to all other responders. The evaluation criterion is whether the information is disseminated *reliably.* To this end, WIISARD employs a peer-to-peer architecture in which information is cached aggressively and disseminated using the gossip-based communication protocol described in Section 4.3.3. All WIISARD devices operate in 802.11 ad-hoc mode.

This architecture has three notable features. First, WIISARD does not require the deployment of additional communication infrastructure: the devices carried by responders should be sufficient to support timely and reliable communication. This decision is motivated by our experiences with the previous version of WIISARD, which required the deployment of mesh nodes to provide coverage prior to the start of drill exercises. The deployment effort was onerous, requiring more than an hour to deploy, configure, and test the mesh nodes. Even worse, once responders and equipment arrived on-scene, the network properties changed dramatically, requiring us to move the mesh nodes. As it took time to notice the network failures and move the nodes, failures could be prolonged. Obviously, such a deployment process would not be viable in real emergencies. It is important to note that WIISARD can take advantage of additional infrastructure to improve connectivity, however, *infrastructure nodes provide a value-added service* rather than being required to support communication.

Second, we opt for a simple broadcast primitive coupled with aggressive caching rather than the more complex communication primitives adopted in other systems (e.g., [10,20, 21,24]). This approach is justified by the fact that a responder may need to access to the information of any patient at any time. A benefit of caching all data is that it facilitates data dissemination through data muling as responders move about the scene. We hypothesize (and verify in Section 5) that data muling plays a significant role in data dissemination. It is important to understand that caching all data is a viable option because WIISARD is a *low-data rate application.* WIISARD focuses on managing patient records along with ambulance and hospital information, which are generated through manual input. Moreover, even our most resource-constrained devices – the mobile phones – can store all the generated data.

Finally, WIISARD minimizes the impact of network partitions and mobility by using a gossip-based communication protocol. The gossip-based protocol relies on *local* communication to disseminate data. This has important advantages over traditional routing protocols that require the construction and maintenance of end-to-end multi-hop routes that are subject to significant temporal dynamics.

4.2 Hardware Components

WIISARD uses hardware that is heterogeneous across three relevant dimensions: display size, operating system, and packet transmission power.

Triage device: The triage device is a Nokia N900 mobile phone that runs Maemo OS – a derivative of the Linux operating system. The small form factor and weight make the triage devices a good choice for the mobile entry and medical teams. The Nokia N900 uses a high-end OMAP 3430 ARM Cortex A8, has 32GB flash storage, and Wi-Fi card with a transmission power of 10mW. The GPS available on Nokia N900 phones is used to track the locations of responders.

Mid-tier/command center: The mid-tier device is a TabletPC while the command center is a large-screen laptop. Their larger touch screens allow for complex user interfaces to manage ambulance and hospital resources. Both devices run Windows XP operating system, are resource rich, and transmit packets at 100mW.

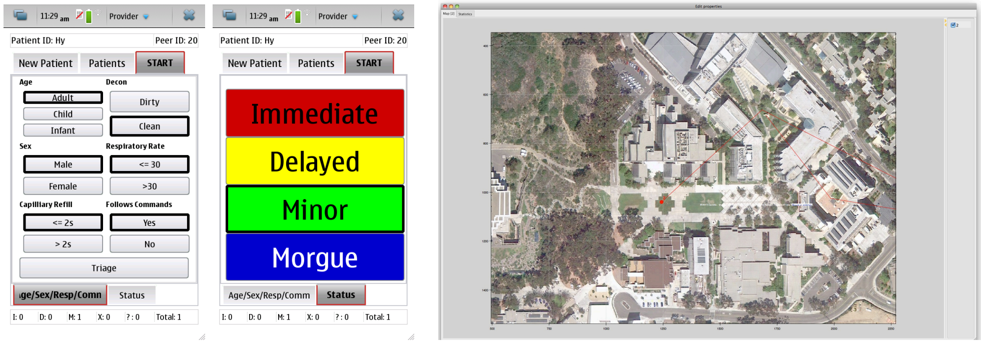

Patient ID: Hy Peer ID: 20

I: 0 D: 0 M: 1 X: 0 ? : 0 Total: 1

(a) START Triage

(b) Command Center

Figure 1: START Triage and Command Center components

Mesh boxes: Mesh boxes may be deployed to improve network connectivity. Each mesh box contains an ALIX x86 board with a 500 MHz AMD CPU and 256 MB RAM, with a 4 GB compact flash card acting as hard drive. Each board also has two mini PCI 802.11abg cards and transmits packets at 100mW. The mesh boxes run Linux Ubuntu 10.04 LTS.

Patient tracking: Traditionally, victims are tracked during emergency responses using paper triage tags that capture the severity of their injuries. To facilitate electronic tracking, we evaluated bar code and RFID technologies. We opted for RFID technology due to the difficulties of using barcodes during sunny days. RFID tags are taped onto the traditional paper triage tags. A NephSystem N330 RFID reader is used to read the short-range RFID tags. The reader is connected to any WIISARD device using Bluetooth.

4.3 Software Components

WIISARD is a portable system that supports both ARM and Intel processors and may be deployed on multiple operating systems including Linux, Windows XP, and Mac OS X. The majority of WIISARD is implemented in Python to ensure cross-platform portability. Only the WIISARD communication protocol is implemented in C++, due to performance considerations. WIISARD has three key components: user interfaces, object storage, and networking.

4.3.1 User Interfaces

The WIISARD user interfaces are developed using the QT toolkit that is accessed via PyQt bindings. The interface design is the result of multiple iterations and interactions with first responders. Particular attention was given to adapting each interface to match the screen size of each device. WIISARD provides the following user interfaces:

Triage: The triage user interface implements the Simple TriAge Rapid Treatment (START) protocol [23] (see Figure 1(a)). The triage of a victim starts by scanning the RFID tag located on their triage tag. The RFID reader transmits the tag ID to the mobile phone over Bluetooth. The START triage interface is designed to enable fast triage by minimizing user input. Typically, a victim is triaged within a minute. Victims may be retriaged when necessary.

Transport: A transport supervisor is responsible for assigning patients to ambulances and designating destination hospitals. This decision is usually based on the nature of victim injuries, need for specialized hospital facilities (e.g., burn center), and hospital bed availability. Traditionally, this information is maintained using paper worksheets. WIISARD improves this process by implementing a transport interface that allows the officer to quickly identify a victim and assign them to ambulances and hospitals. Victims are identified either based on their injury status or by scanning their RFID tag. This enables responders to quickly bring up victim information. The more complex transport user interface is effectively supported by the larger screen size of the tablet devices compared to that of the mobile phones.

Command Center: The command center is capable of accessing all patient information, creates summaries about the progress of the response, and displays the locations of first responders and victims (see Figure 1(b)). WIISARD uses the GPS capabilities of mobile phones to track not only the location of the first responders that carry the mobile phones but also the location of the victims. The location of victims is inferred based on the location of the providers: when a provider scans a victim's RFID, the victim's location is updated based on the provider's current location.

4.3.2 Object Storage

The data generated by responders is modeled as objects that are persisted on disk using SqlLite. Each object has a globally unique identifier. In the case of patient data, their RFID tag is used as the unique identifier; the identifiers of the other objects are randomly generated Universally Unique IDentifiers (see RFC 4122).

WIISARD distinguishes between different versions of the same object through time stamps. Obviously, this requires all peers to be time synchronized. WIISARD achieves time synchronization either through NTP or GPS time. Version conflicts are resolved by keeping the object with the latest timestamp. While this does not eliminate the potential for version conflicts, such conflicts would seldom occur in practice because of two factors. First, it is unlikely for first responders at different locations to update the same record simultaneously. In fact, most of the time the responders are within physical proximity of the victim to update the records, as they need to scan their triage tag. Second, since data is entered manually, small clock synchronization errors have minimal impact on distinguishing between recent and old version of an object.

The object storage component implements two key functions. First, the object storage mediates the interactions between the network component and the user interface components through a standard model-view-control design pattern. A user interface component registers callbacks through which it is notified of object updates. Second, WIISARD is designed to recover from application crashes. WIISARD implements the following check-point policy: the data generated locally is committed to the database as soon as possible to ensure its persistence, while data received from other nodes is buffered before being committed to disk. The policy does not introduce high overhead as the users create or modify objects infrequently.

4.3.3 WIISARD Communication Protocol

The initial prototype of WIISARD used a client-server architecture and an AODV-based routing protocol [14, 24]. This initial prototype suffered from inconsistent network reliability across deployments. The poor reliability is partly explained by the need to construct and maintain multi-hop paths in a highly dynamic network environment. Moreover, network partitions disconnect the clients and server, preventing communication.

A pragmatic approach to improving reliability is to forgo the client-server architecture that requires the use of end-to-end routes and opt for a peer-to-peer architecture in which data is disseminated through a gossip-based protocol that relies on local communication. This approach is a good fit with the communication requirements of emergency responses, where information must be shared across all responders. A gossip protocol works by having each peer "gossip" the information it receives until all peers within the network share this information. The fundamental challenge of gossip protocols is to avoid the *broadcast-storm problem*: a peer requests a piece of information and its neighbors rush to send it, resulting in packet collisions and high overhead.

To meet this challenge, we developed the **W**IISARD **C**ommunication **P**rotocol (WCP), a reliable and efficient gossip protocol. A preliminary evaluation of WCP is presented in [14]. WCP divides the user-generated data into Blocks such that each Block fits in a packet. To improve response time and minimize memory utilization, WCP uses a caching scheme to maintain recently referenced blocks in memory while committing the rest to the object storage, as described in Section 4.3.2.

WCP works as follows. Peer n divides its local time into periods of length P and each period has W equally sized slots. Note that each peer operates independently without requiring time synchronization. In the beginning of each period, n transmits a Beacon that summarizes the blocks stored in its local storage component as version vectors [32]. Upon receiving a Beacon from a neighbor m, n inspects its object store to determine if it has any Blocks that m does not. Peer n selects random slots to transmit the Blocks that m is missing within the W slots following the reception of m's Beacon. Accordingly, in response to a Beacon, up to W Blocks may be transmitted. By picking the slot at random, the source of the packet is randomized over multiple requests. This avoids selecting a peer with poor link quality from transmitting in multiple rounds.

WCP uses two optimizations to reduce the number of duplicate packets. First, a peer n does not transmit the same Block more than once within W slots. This reduces the overhead when multiple neighbors are missing the same block. Peer n may fulfill many outstanding requests using a single transmission by taking advantage of the broadcast nature of wireless communication. Second, peer n cancels the transmission of a Block when it overhears the Block multiple times (twice in our deployment). Such an optimization is particularly effective in dense connectivity graphs in which many peers may receive the broadcast Block. This optimization was initially proposed as part of Trickle [25].

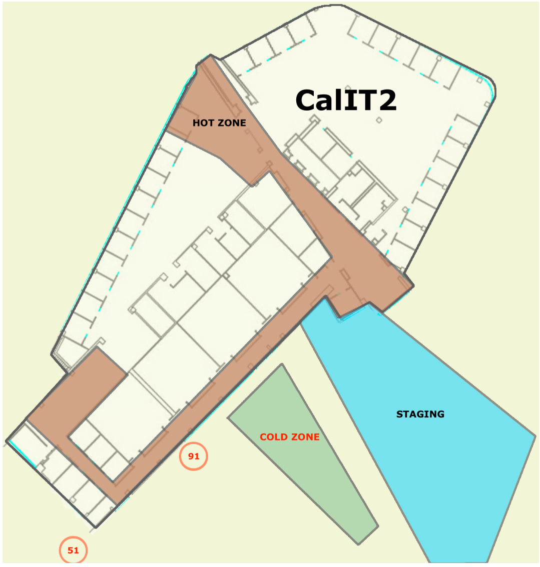

Figure 2: Physical layout of the drill

Variable	Statistic
Responders	16 responders / 3 supervisors
Victims	41 victims
Deployed devices	15 triage devices 3 mid-tiers 1 command center 2 mesh boxes
Phases of the drill	staging (11:00:00 – 11:06:00) rescue (11:06:00 – 11:32:00) treatment (11:32:00 – 11:37:00)

Table 1: Deployment statistics

WCP makes minimal assumptions about network connectivity. WCP relies on *soft state* that is refreshed periodically through Beacons. In fact, the only assumption WCP makes is that network connectivity remains stable within a beacon period P. The beacon period represents a trade-off between the protocol overhead and the assumed stability of the network. In our deployment, P is set to 5 seconds. Moreover, since WCP caches all data, it can also disseminate data through data muling as opportunistic connections are established with other peers. The use of local communication, soft state, and data muling enables WCP to achieve a network reliability of 98% during the drill exercise presented in Section 5. This is in spite of a highly dynamic environment characterized by variable link quality, mobile users, and network partitions.

5. DRILL EXERCISE

We deployed WIISARD as part of a drill organized at University of California, San Diego (UCSD) on August 10th, 2011. Drill exercises are a common and effective method of training and evaluating the readiness of emergency responders. The WIISARD team had previously participated in several drills administered by San Diego County emergency agencies, and our emergency doctors had been part of the

County team that designs and stages drills, enabling us to create a realistic drill.

The drill scenario called for a major earthquake that injured and trapped multiple victims on the fifth floor of Atkinson Hall. The mission of the first responders was to assess the criticality of injuries sustained by victims, remove them from the building, and provide necessary medical treatment.

The exercise involved 19 responders and 41 victims. The first responders were fire fighters from stations located in the vicinity of UCSD, plus two ambulance crews from Pacific Ambulance. The victims were volunteers mainly from the San Diego and UCSD chapters of the Community Emergency Response Team (CERT). Victims were provided with scenario cards that described the extent of their injuries. Scenarios were designed by our medical collaborators to cover a range of injuries consistent with those that may be sustained in earthquakes. To increase the realism of the drill, moulage was applied and victims were encouraged to act out their injuries: scream, act disoriented, complain, and ask for preferential treatment. Moreover, some of the office furniture on the fifth-floor of Atkinson Hall was positioned to simulate the impact of an earthquake.

Figure 2 shows the geographic layout of the drill. The hot zone is located on the fifth floor of Atkinson Hall. The cold zone is located in front of the building. The emergency response proceeded in three phases: *staging*, *rescue*, and *treatment*. The duration of each phase is captured in Table 1. During the *staging phase*, the first responders arrive on scene, and establish command and control. We instrumented the entry team and the transport and medical supervisors. During the analysis we will refer to the members of the entry team as ***responders*** and to the transport and medical supervisors as ***supervisors***. In the *rescue phase*, the entry team moved in the building to locate, triage, and evacuate the victims to the cold zone. Meanwhile, most of the supervisors remained in the cold zone. Once all victims were transported to the cold zone, the *treatment phase* started. During this phase the entry/medical teams retriaged victims and provided additional details regarding their injuries.

The deployment included all components of the WIISARD system: 16 triage devices, three mid-tiers, and a Command Center (see Table 1). As our goal was to create a realistic deployment scenario in which minimal infrastructure is used, we only include two mesh boxes. The mesh boxes were deployed outdoors as indicated by the circled numbers in Figure 2. The mesh boxes had minimal impact on the observed results: they improved connectivity during treatment. However, the boxes had no impact either during staging when they were off or during rescue when they received no data from within the building. We supplied the first responders with triage tags augmented to each include an RFID tag. As previously discussed, the triage tags are usually placed around the victims' necks to track them during the response.

To create a detailed communication record, each peer recorded the packets it transmitted and received. Since all peers communicated via a shared multicast address, each peer records the messages transmitted by all other peers within its communication range. The minimum transmission rate is 0.2 packet per second as WCP transmits beacons every 5 seconds. The WIISARD devices were time synchronized using NTP in the beginning of the drill and whenever connections to the NTP server (deployed on the command center) were established. We performed a retrospective analysis to verify that the peers were time synchronized by comparing the timestamps when different peers received the same packet. The analysis revealed a time synchronization error below a second.

6. NETWORK PROPERTIES

In this section we characterize the network properties that emerge during emergency responses, based on in-situ measurements from the drill exercise. Specifically, our study focuses on the following questions:

1. What are the underlying network properties of emergency responses, including distribution of link failures and density/diameter of connectivity graph?
2. How do network properties vary across the phases of the drill and with the roles responders have in the drill?
3. What is the impact of mobility on network properties?

Answering these questions provides a sound basis for understanding the challenges of reliable communication in emergency response and will guide the selection appropriate network architectures and protocols to overcome these challenges. Our analysis focuses on understanding the challenges of reliable communication when minimal infrastructure is deployed. We characterize both the properties of links and those of the connectivity graph. A common thread in our analysis is that the emergent network properties are the result of the emergency response workflow and team structure.

6.1 Link Properties

We investigate the reliability of links established among responders during the drill. We start by considering whether link properties change across the phases the drill or with the responder's roles in the drill. Data analysis indicates that the rescue phase posed the most significant challenges to reliable communication. Next, we characterize the long-term and the short-term variations in link quality during the rescue phase. The analysis of link properties at different time scales is motivated by the fact that different mechanisms are employed to handle short-term and long-term variations in link quality.

Mobility is a primary cause of long-term link variability: links are established and broken as the responders rescue victims. To quantify the long-term variations in links, we define *contacts* between pairs of peers. A contact (n, m) means that a peer n is in contact with peer m, while the time difference between consecutive packets that n receives from m is less than a minute. We use *contact length* and *inter-contact time* to assess the temporal properties of contacts. The *contact length* and *inter-contact time* are computed based on both beacon and data packets.

Of interest are the short-term variations in link quality during contacts. Many factors contribute to short-term variations in link quality, including interference, changes in antenna orientation, and movement over short distances. We quantify short-term link quality by measuring the packet reception rates (PRRs) during peer contacts. PRR is computed based on the beacons transmitted by WCP.

6.1.1 Link Properties Vary between Drill Phases

Figure 3(a) plots the number of distinct contacts of responders during a one-minute window. The background

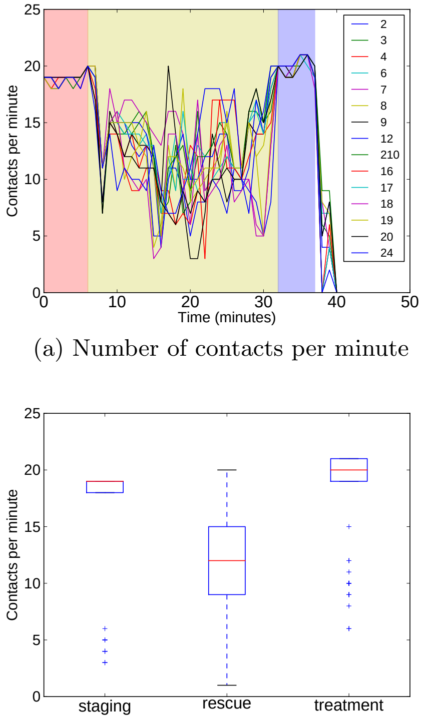

(a) Number of contacts per minute

(b) Distribution of contacts for responders

Figure 3: Number of contacts during the drill

color indicates the phase of the drill: staging, rescue, or treatment. During the staging and treatment phases, responders had a similar average number of contacts. According to Figure 3(b), the median number of contacts during staging and treatment was 19 and 20, respectively. The staging phase had little variability: the range was 18 – 19 contacts per minute. This is the result of the close proximity of responders as they planned their response. During treatment we observe a wider range of 6 – 21 contacts per minute. The higher maximum number of contacts was due to the deployment of two mesh boxes at the end of the staging phase. The increased variability in the number of contacts may be attributed to the larger area in which the responders were distributed and increased mobility as responders re-triaged and treated patients.

The number of contacts during the rescue phase (median 12, range 1 – 21) differs significantly from either the staging or treatment phases. The lower median value and higher variability are the compounded result of the indoor environment where walls significantly attenuated links as well as the increased mobility of the responders during the rescue phase. This result indicates that network properties change significantly among the different phases of the drill.

Result: *The long-term link properties vary across the phases of the drill. The rescue phase exhibits high variability in number of contacts per minute due to mobility.*

6.1.2 Link Properties depend on Responder's Roles

An emergency response system must disseminate data both among responders and between responders and supervisors. We define the link groups – `rr` and `rs` – to include all pairs of links that may be established among responders (`rr`) and between responders and supervisors (`rs`), respectively. To determine whether link properties depend on the responder's role in the drill, we compare the properties of links in `rr` and `rs`.

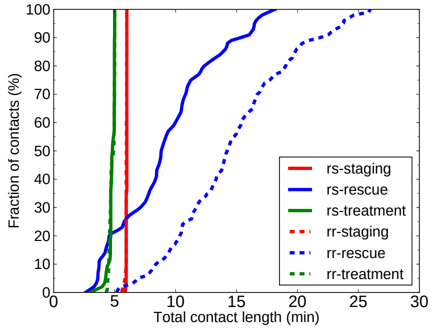

Figure 4: CDF of total contact lengths among rescuers (`rr`) and between rescuers and supervisors (`rs`) during each phase of the drill

To this end, we define the *total contact length* of peers (n, m) as the sum of their contact lengths during each phase. The CDF of the total contact lengths during each phase is plotted in Figure 4.

During staging and treatment, there were small differences in the distribution of total contact lengths between the `rr` and `rs` groups. More importantly, a large fraction of the links was established for more than 90% of their respective phase durations. The sharp increase in the CDF for `rr` and `rs` during staging and treatment indicate the end of the phase after 5 and 6 minutes, respectively. This indicates stable connectivity among responders and between responders and supervisors.

During the rescue phase, there was good connectivity among responders. All contacts among responders (`rr`) had a total contact length of at least 5 minutes. Moreover, half of the contacts in `rr` had total contact lengths longer than 14 minutes. In sharp contrast, the connectivity between responders and supervisors was poor: 22% of the contacts in `rs` had a total contact length less than 5 minutes (there were no such short contacts in `rr`). Even worse, there is no contact between responders and supervisors (`rs` group) that has a total length longer than 17 minutes (i.e., 69% of the rescue phase) indicating that the two groups were disconnected for part of the rescue phase.

This data indicates that there was good connectivity both in the `rr` and `rs` groups during treatment and staging. However, the responders and supervisors (`rs`) were disconnected for part of the rescue phase. A likely explanation is that the supervisors stationed themselves in the cold zone, while the rescuers that were part of the entry team proceeded into the hot zone to rescue the victims. This shows that it may be more difficult to disseminate data from responders to their supervisors than it is among the responders.

Result: *The long-term link properties depend on the responder's roles in the drill. During the rescue phase, the connectivity between responders and supervisors may be intermittent as responders enter the hot zone and supervisors remain in the cold zone.*

6.1.3 *Long-term Variations in Link Quality*

Next, we characterize the long-term variations of links as measured by the contact metrics. The presented data shows that the rescue phase poses the most significant challenges to reliable communication. Accordingly, we will focus the remainder of our analysis on the network properties observed during this phase.

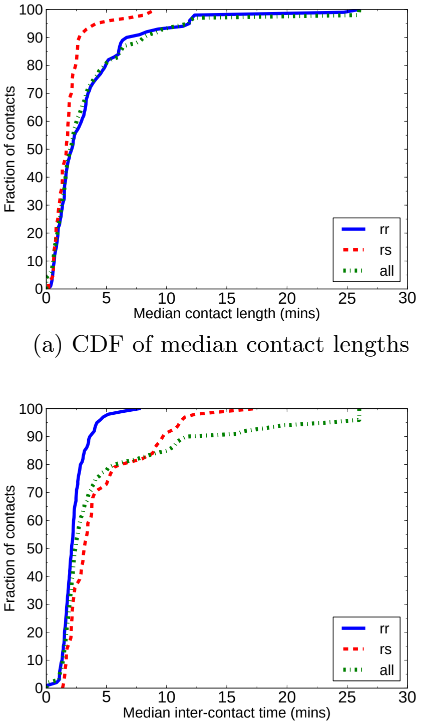

(a) CDF of median contact lengths

(b) CDF of median inter-contact times

Figure 5: CDF of median contract lengths and median inter-contact during rescue phase

Figure 5(a) plots the CDF of median contact-lengths for the `rr` and `rs` groups during rescue. Half of the links had a median contact-length of 1.94 minutes. The short median contact length is indicative of a dynamic network in which link failures are common. Moreover, the contacts among responders were longer than those between responders and supervisors. For example, 90% of the contacts in `rs` were shorter than 2.75 minutes, compared to only 58% of the contacts in `rr`. Another interesting property of the distribution is its long tail, particularly for the `rr` group of contacts. The long tail is due to a small subset of contacts that were stable for prolonged periods of time. In fact, a small fraction of the links was stable for the entire 26 minutes of the rescue. This is consistent with our empirical observation that responders tend to work in small groups. It also stands in contrast to prior empirical DTN studies in which this behavior was attributed to popular locations rather than teamwork.

Figure 5(b) plots the CDF of inter-contact lengths. Half of the contacts had a median inter-contact length of about 2.15 and 3.15 minutes for the `rr` and `rs` groups, respectively. Overall, contacts between responders were reestablished significantly faster than those occurring between responders and supervisors. In fact, 90% of the contacts in `rr` were reestablished within 2.8 minutes compared to 6.85 minutes for contacts in `rs`. As previously discussed, this is the result of the supervisors typically being located in the cold zone while the responders entered the hot zone.

Result: *Contacts between responders are typically short, however, they are often reestablished within minutes. Moreover, a subset of responders establishes contacts that are stable for prolonged periods of time.*

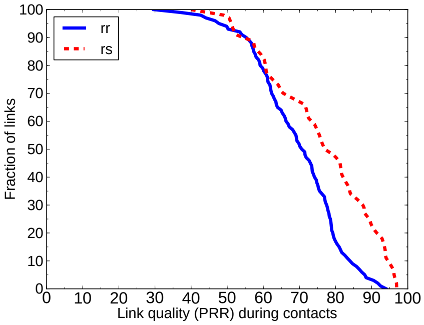

Figure 6: Complementary CDF of link quality as measured by PRR while peers are in contact

6.1.4 *Short-term Variations in Link Quality*

Figure 6 plots the complementary CDF of the observed PRR while peers are in contact. This result captures the variability in link quality that can be attributed to environmental factors including wall attenuation, interference, and changes in antenna orientation due to body movement. The figure indicates that once a contact was established, the link quality tended to be relatively high. For example, the median PRR of links in `rr` and `rs` exceeded 70% and 77%, respectively. A consequence of the good PPR observed while contacts were established is that retransmissions were effective in combating short-term variations in link quality.

Result: *The short-term link quality while contacts are established is high.*

6.2 Connectivity Properties

This section investigates the connectivity of the responders during the rescue phase. The analysis focuses on the key factors that affect the reliability and latency of data dissemination including: the prevalence of network partitions, the density and diameter of the connectivity graph, and the potential of using data muling to disseminate data. The connectivity graph was computed by merging the time synchronized logs of peers.

6.2.1 *Network Partitions*

For the majority of the rescue phase, all fifteen responders were located inside Atkinson Hall triaging patients. Due to the small physical area that was in play (see Figure 2), we expected that the network would remain connected as responders moved on the floor. Figure 7(a) plots the number of connected components in the communication graph. While for much of the rescue phase the graph was connected (i.e., it had a single connected component), to our surprise, the network was partitioned for 26.3% of the phase. Three factors contribute to this result. First, the responders operated indoors where walls limit signal propagation. Second, the transmission power of phones was 10 mW (compared to 100mW for the other devices). Finally, responders tended to operate in groups that were not distributed uniformly throughout the building.

Partitions can have a profound impact on network performance by preventing end-to-end paths to be established. Figure 7(b) plots the fraction of responder pairs for which there is no path at a time instant. During the rescue phase, 67% of responder pairs were affected by partitions. Thus, traditional routing protocols that require continuous connectivity may not be suitable for emergency responses.

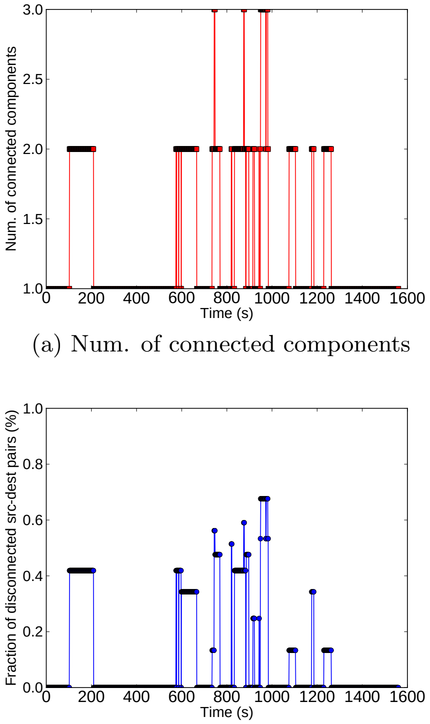

(a) Num. of connected components

(b) Fraction of disconnected peers

Figure 7: Impact of network partitions

Result: *Network partitions are common and prevent many pairs of peers from communicating.*

6.2.2 Path Lengths and Density of the Connectivity Graph

Figure 8(a) plots the median and maximum path length (i.e., diameter) of all paths between responders during the rescue phase. The distribution of these metrics over time is plotted in Figure 8(b). For 73% of the rescue phase, the diameter of the graph was at most three while the median path length was one. During this time, the connectivity graph was dense and had a low diameter. Such graphs indicate low communication latencies and effective data dissemination, as a small number of broadcasts were sufficient to relay data to all peers. However, for the remainder of the time, the communication graph was rather sparse: the diameter was as large as seven with a median path length as high as three. These results indicate graphs with well-connected components (small median path length) that were connected through a few links (large diameters). The diversity of the graph structures observed during the rescue phase underline the need to develop communication algorithms that operate well in both dense and sparse communication graphs.

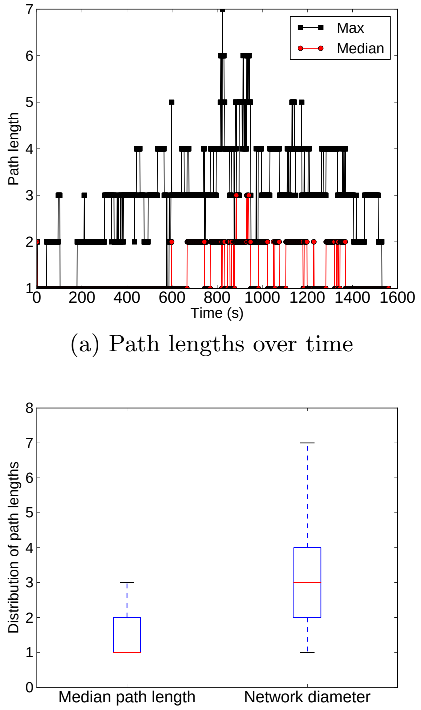

(a) Path lengths over time

(b) Distribution of path lengths over time

Figure 8: Path lengths statistics

Result: *Reliable communication must be supported in both dense and sparse graphs.*

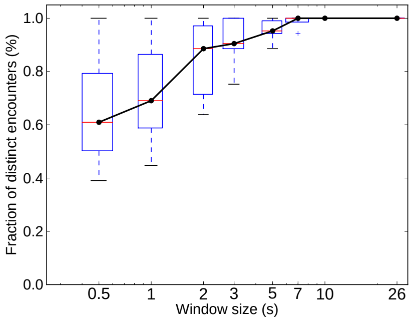

Figure 9: Fraction of encountered nodes over for different windows sizes during rescue phase

6.2.3 The Benefits of Mobility

DTN techniques enable protocols such as WCP to take advantage of mobility to improve data dissemination. To quantify the potential of delivering data through data muling, for each responder we compute the fraction of other responders they encountered during windows of different sizes. Figure 9 plots the distribution of encounters for each responder for a given window size. As the window size increases, a responder tends to encounter an increasing number of the other responders. The number of encountered peers increases at an exponential rate over time (note that the x-axis is logarithmic). This result is captured both in the increasing median values as well as the rapidly decreasing variance. In fact, within seven minutes, a responder encountered all other responders. This shows that if we are willing to tolerate a transmission latency of seven minutes, data among responders can be disseminated solely through data muling.

Previous studies [15] that considered the encounter patterns of nodes within WLAN traces of users show that it is unlikely for a user to meet a large fraction of the other users. The fundamental difference between the prior studies and our results is that we are specifically looking at the encounter patterns of first responders that work cooperatively to rescue victims, rather than considering populations of unrelated users. This indicates that disconnected routing protocols should be a good choice for delivering data among first responders during emergency responses.

Result: *Data dissemination may be improved through data muling as mobile responders rescue victims.*

7. NETWORK PERFORMANCE OF WIISARD

The previous section characterizes the link and network properties that are observed during drills. These results highlight the significant challenges that must be overcome to support reliable communication during emergency responses. As discussed in Section 4.3.3, WCP addresses these challenges through a simple, yet robust gossip-based protocol. In contrast to traditional routing protocols that construct and maintain end-to-end routes, WCP relies on local message exchanges to tolerate variations in link quality and takes advantage of data muling to disseminate data.

During the deployment, victim records were generated during triage/retriage and location records were generated when valid GPS coordinates were acquired. Figure 7 plots the number of generated reports by each responder that used a mobile phone. As responders spent a similar time outdoors, they transmit similar numbers of location reports. In contrast, the number of victim records varied significantly indicating that responders triaged an uneven number of victims. Figure 10(b) plots the total number of reports transmitted (including retransmissions) by all peers during each minute of the drill. On average, a total of 13 packets per second including a single record were transmitted during the drill. However, the sending rate varied significantly observing a maximum aggregated transmission rate of 33 packets per second. The increases in workload are correlated with WIISARD recovering from network partitions. Overall, the data indicates that the system is below its network capacity, indicating the potential to scale to larger numbers of responders. Moreover, the sending rate may be significantly reduced by transmitting multiple records in a single packet, since the typical size of a record is about 200 bytes, far below the maximum payload of 802.11 packets.

The goal of WIISARD is to deliver the data reports generated by a node to all other nodes. We say that a report was delivered reliably, if all nodes in the network receive a copy of the data report. Note that even if a single node does not receive the report, we count this as a failure of the protocol. Even with this strict requirement, WCP performed well: WCP delivered 98.25% of the generated data reports to all nodes. Most of the dropped data reports occurred towards the end of the trace, when the most recently captured patient data might not have had time to propagate through the network like data captured earlier.

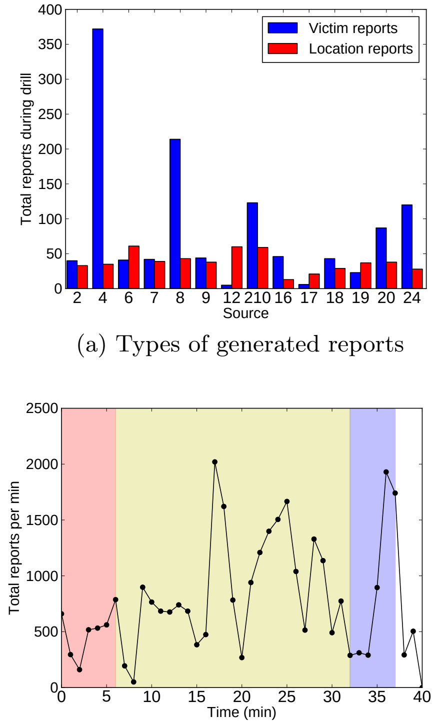

(a) Types of generated reports

(b) Total transmitted reports over time

Figure 10: Characterization of transmission rate

Figure 11(a) plots the reliability of delivering the data reports created by each node. In accordance to the overall reliability of 98.25%, most of the data reports generated by a node were successfully disseminated. The clear exception is peer 17 for which only 80% of the data reports it created were disseminated to all other nodes. In part, this result is skewed due to the small number of data reports generated by peer 17. We note that no data reports were generated on the mid-tier or command center; these devices were used primarily to view victim statistics.

Figure 11(b) plots the CDF of dissemination latency for the subset of data reports that were delivered to all peers. The median and 90th percentile of the distribution were 30 seconds and 2.2 minutes, respectively. However, the distribution also has a long tail: the worst-case latency was 9.1 minutes. This indicates that while most of the time data was delivered within 2.2 to most responders, due to network partitions, some responders experienced prolonged delays.

8. DISCUSSION

Impact of Emergency Response Workflow: The presented study highlights the profound impact of the emergency response workflow on the mobility patterns and network properties observed during the drill. As a result, emergency response differs from the previously studied DTN applications in the following key aspects. First, link properties

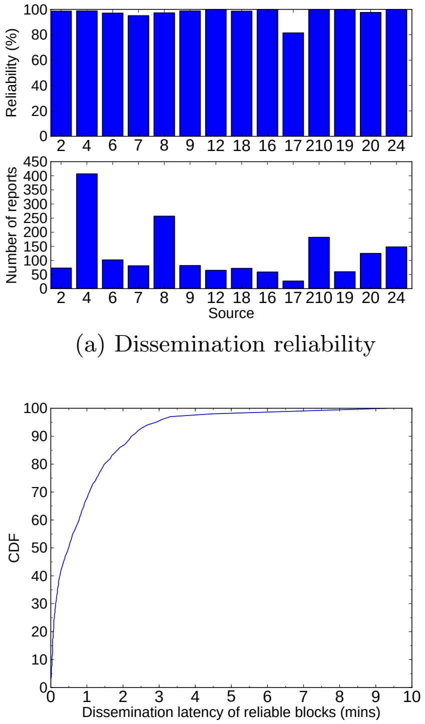

(a) Dissemination reliability

(b) CDF of dissemination latency

Figure 11: Dissemination reliability and latency

vary across the phases of the drill and with the roles of the responders. In contrast, prior DTN studies often assume that the mobility patterns of users are statistically similar and that network properties are stationary [8,11,34,35]. Second, the team structure – rather than popular locations [8,34] or social relationships [15,31] – are responsible for the long tail of the cumulative distribution of median contact lengths. Finally, the emergency response workflow offers unexpected opportunities for disseminating data through data muling: when (1) teams physically transfer patients from one to another and (2) as teams enter/exit the hot zone in which access is managed through a single or few entry points. This shows that emergency response has a more complex network properties that can be understood starting from the emergency response workflow.

Generalizability: Emergency response belongs to a larger class of applications that involve mobile entities that perform cooperative tasks according to well-established workflows. A typical workflow may dictate team structure, physical movement, and best practices for operation. In the case of emergency response, these aspects are defined as part of the ICS (as discussed in Section 3). Other applications that pertain to this class include military operations, participatory sensing tasks (e.g., surveys of earthquake-damaged buildings), or the operation of robotic teams/swarms (e.g., crop pollination [16]). Next, we discuss the impact of the empirical properties on emergency response. However, we expect that some of the observed challenges and associated solutions may be applicable to this application class.

Networking Architecture: The presented results provide a strong case against the use of traditional mesh routing protocols in emergency response systems. First, mesh routing requires end-to-end routes to be maintained in spite of network dynamics. Our empirical data clearly indicates the presence of significant network dynamics both within and across the phases of the drill. Second, mesh routing protocols have not been designed to cope with network partitions, which are common even in the relatively confined area in which the drill occurred. These two factors explain the poor performance that we observed with the initial WIISARD system (10% data reliability) and are consistent with those obtained as part of the Code Blue Project [26] (20% data reliability according to [12]). In contrast, WIISARD achieves 98.25% data delivery during the drill. This shows the superiority of DTN techniques over mesh routing protocols for emergency response systems.

State Consistency: Brewer's CAP theorem [5] states that a distributed system cannot simultaneously provide consistency, availability, and tolerate partitions. WIISARD is an "always available" system: it provides responders with the most up-to-date information available. Therefore, in systems such as WIISARD, there is a fundamental tradeoff between consistency and tolerance to network partitions. WIISARD supports an eventual consistency model: WIISARD cannot guarantee consistency when the network is partitioned, however, WIISARD will automatically update its records to their latest versions as partitions merge. In WIISARD, inconsistent state can give rise to a situation where a responder or incident commander has an incorrect view of the status of a patient, such as whether the patient has been assigned a triage level.

An important property of WIISARD is that it requires different degrees of consistency based on the tasks and roles of providers. The primary function of WIISARD is to support cooperative triage and treatment of patients by the members of a response team. It is essential for these members to have consistent views of patient records. In practice, it is easy to achieve consistency among members of the same response team since they are usually within physical proximity and, as a result, they are part of the same connected graph component. Since according to the emergency response workflow patients are physically handed off from one response team to another, this ensures the presence of opportunities of synchronizing patient records between teams through data muling. An additional measure that may be taken to ensure consistency of patient records between teams is to provide each patient with a device that stores and forwards their patient records.

The secondary function of WIISARD is to provide incident commanders with statistics regarding patient counts and the distribution of injury severity, which requires less strict consistency guarantees. In our drill exercise we observed maximum delays of 9.1 minutes to deliver data to incident commanders. In its current implementation WIISARD does not provide differentiated service based on the consistency guarantees required by responders. We expect that such traffic differentiation could improve the performance of WIISARD.

Limiting Network Partitions: There are a number of opportunities to improve the performance of WIISARD. A non-technological approach to minimizing the observed latencies is to modify the response workflow to start the evacuation of victims earlier. This will ensure that data can be delivered from the hot zone to the cold zone through data muling. Technological alternatives focus on minimizing the likelihood of observing network partitions. A pragmatic approach is to deploy additional infrastructure to bridge the cold and hot zones since our study indicates that DTN techniques can already provide reliable communication within each zone. Alternatively, a long-range wireless technology, carried by responders or mounted on their vehicles, may be employed to a similar effect. However, as long-range wireless solutions usually have low bandwidth, it is important to prioritize patient information. The development of tools for network deployment and the evaluation of the effectiveness of combining long-range and short-range wireless technology remain research questions in emergency response systems.

Testing and Simulations: A fundamental challenge to developing robust emergency response system is testing: there are only limited opportunities to deploy a system as part of drill exercises to test their performance. Months of software testing typically go into preparing for a drill exercise to ensure correct operation stifling innovation in this area. For example, without an understanding of mobility patterns of responders, we had to opt for a gossip-based protocol that does not account for user mobility. The presented results highlight that the mobility patterns of responders tend to be structured. We plan on developing protocols that take advantage of these properties to improve delivery performance, particularly focusing on reducing the latency of packet delivery. Moreover, the insights presented here, combined with the collected data can be the basis for developing a realistic simulation environment for evaluating emergency response systems.

9. CONCLUSIONS

This paper presented the design, deployment, and empirical evaluation of WIISARD during a drill exercise. The drill exercise is designed to assess the feasibility of achieving reliable communication with limited communication infrastructure. The drill involved 19 instrumented first responders that triaged 41 simulated victims. Our work makes the following key contributions.

1. We provide a detailed characterization of the network properties that emerge during an emergency response. We show that network properties are highly variable and depend on both the drill phase and the responder's role during the emergency response. The contacts between responders tend to be short, however, they are often reestablished within minutes. The quality of links during contacts tends to be high. The connectivity graph usually has small diameter and is dense. However, due to limited infrastructure, at times it becomes sparse and even partitioned. Network partitions can be effectively mitigated by taking advantage of the frequent contacts between responders, which are an important characteristic of the emergency response workflow.

2. In spite of these challenges, WIISARD achieves a reliability of 98% during the drill. This shows the effectiveness of using gossip-based communication and taking advantage of data muling.

3. The empirical results also point towards several ways in which emergency response systems may be improved, including support supporting differentiated consistency guarantees, leveraging the structured motion patterns of responders to reduce communication latencies, and the development of a sound testing environment for such systems. The presented study highlights the profound impact of the emergency response workflow on the mobility patterns and network properties observed during the drill.

Acknowledgments

We express profound thanks to the first responders that participated in our emergency response drill and allowed us to work with them to evaluate WIISARD. We also thank the other members of the WIISARD team who played critical roles in preparation for exercises and administration of the project. Finally, we would like to thank our anonymous reviewers for their feedback and our shepherd Matt Welsh.

This work was supported by NIH/NLM grant number R01-LM009522 and by NSF through grants CNS-1144664 (NeTS) and CNS-1144757 (NeTS).

10. REFERENCES

[1] Emergency Patient Tracking System (EPTS). `http://www.raytheon.com/capabilities/products/epts/`.

[2] G. A. Bigley and K. H. Roberts. The incident command system: High-reliability organizing for complex and volatile task environments. *The Academy of Management Journal*, Vol. 44(No. 6):pp. 1281–1299, 2001.

[3] D. Bradler, B. Schiller, and E. Aitenbichler. Towards a distributed crisis response communication system. In *Proceedings of ISCRAM*, 2009.

[4] B. Braunstein, T. Trimble, R. Mishra, B. S. Manoj, R. Rao, and L. Lenert. Feasibility of using distributed wireless mesh networks for medical emergency response. In *Proceedings of AMIA*, 2006.

[5] E. A. Brewer. Towards robust distributed systems (Invited Talk). Principles of Distributed Computing, 2000.

[6] S. W. Brown, W. G. Griswold, B. Demchak, and L. A. Lenert. Middleware for reliable mobile medical workflow support in disaster settings. In *Proceedings of AMIA*, pages 309–313, 2006.

[7] R. Bruno and M. Conti. Opportunistic networking overlays for ICT services in crisis management. *Proceedings of ISCRAM*, 2008.

[8] J. Burgess, B. Gallagher, D. Jensen, and B. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networks. In *Proceedings of INFOCOM*, 2006.

[9] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless communications and mobile computing*, 2(5):483–502, 2002.

[10] R. Carella and S. McGrath. ARTEMIS personal area networks for emergency remote triage and information management. *Proceedings of ISCRAM*, 2006.

[11] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing*, 6:606–620, 2007.

[12] B. Chen, G. Peterson, G. Mainland, and M. Welsh. Livenet: Using passive monitoring to reconstruct

sensor network dynamics. *Distributed Computing in Sensor Systems*, pages 79–98, 2008.

[13] B.-r. Chen, K.-K. Muniswamy-Reddy, and M. Welsh. Ad-hoc multicast routing on resource-limited sensor nodes. In *Proceedings of REALMAN*, pages 87–94, 2006.

[14] O. Chipara, A. N. Plymoth, F. Liu, R. Huang, B. Evans, P. Johansson, R. Rao, and W. G. Griswold. Achieving reliable communication in dynamic emergency responses. In *Proceedings of AMIA*, 2011.

[15] E. Daly and M. Haahr. Social Network Analysis for Information Flow in Disconnected Delay-Tolerant MANETs. *IEEE Transactions on Mobile Computing*, 8(5):606–621, 2009.

[16] K. Dantu, B. Kate, J. Waterman, P. Bailis, and M. Welsh. Programming micro-aerial vehicle swarms with karma. In *Proceedings of SenSys*, Nov. 2011.

[17] R. Dilmaghani and R. Rao. An Ad Hoc Network Infrastructure: Communication and Information Sharing for Emergency Response. In *Proceedings of WIMOB*, 2008.

[18] R. Dilmaghani and R. Rao. A wireless mesh infrastructure deployment with application for emergency scenarios. In *Proceedings of ISCRAM*, 2008.

[19] F. Dong, Y. Hu, M. Tong, and X. Ran. Supporting Emergency Service by Retasking Delay-Tolerant Network Architecture. In *Proceedings of MSN*, 2009.

[20] T. Gao, T. Massey, L. Selavo, D. Crawford, B.-r. Chen, K. Lorincz, V. Shnayder, L. Hauenstein, F. Dabiri, J. Jeng, A. Chanmugam, D. White, M. Sarrafzadeh, and M. Welsh. The Advanced Health and Disaster Aid Network: A Light-Weight Wireless Medical System for Triage. *IEEE Transactions on Biomedical Circuits and Systems*, 1(3):203–216, 2007.

[21] S. George, W. Zhou, H. Chenji, M. Won, Y. O. Lee, A. Pazarloglou, R. Stoleru, and P. Barooah. DistressNet: a wireless ad hoc and sensor network architecture for situation management in disaster response. *IEEE Communications Magazine*, 48(3):128–136, 2010.

[22] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of MobiCom*, 2004.

[23] C. A. Kahn, C. H. Schultz, K. T. Miller, and C. L. Anderson. Does start triage work? an outcomes assessment after a disaster. *Annals of Emergency Medicine*, 54(3):424 – 430.e1, 2009.

[24] L. A. Lenert, D. Kirsh, W. G. Griswold, C. Buono, J. Lyon, R. Rao, and T. C. Chan. Design and evaluation of a wireless electronic health records system for field care in mass casualty settings. *Journal of American Medical Informatics Association*, 18(6):842–852, 2011.

[25] P. Levis, N. Patel, D. Culler, and S. Shenker. Trickle: a self-regulating algorithm for code propagation and maintenance in wireless sensor networks. In *Proceedings of NSDI*, 2004.

[26] K. Lorincz, D. Malan, T. Fulford-Jones, A. Nawoj, A. Clavel, V. Shnayder, G. Mainland, M. Welsh, and S. Moulton. Sensor networks for emergency response: challenges and opportunities. *Pervasive Computing, IEEE*, 3(4):16–23, 2004.

[27] L. McNamara, C. Mascolo, and L. Capra. Media sharing based on colocation prediction in urban transport. In *Proceedings of MobiCom*, 2008.

[28] M. McNett and G. M. Voelker. Access and mobility of wireless PDA users. *ACM SIGMOBILE Mobile Computing and Communications Review*, 9(2):40–55, Apr. 2005.

[29] S. F. Midkiff and C. W. Bostian. Mikobos - a mobile information and communication system for emergency response. In *Proceedings of ISCRAM*.

[30] T. J. Morris, J. Pajak, F. Havlik, J. Kenyon, and D. Calcagni. Battlefield Medical Information System-Tactical (BMIST): The application of mobile computing technologies to support health surveillance in the Department of Defense. *Journal of Telemedicine and E-Health*, 12(4):409–416, 2006.

[31] M. Musolesi and C. Mascolo. Designing mobility models based on social network theory. *ACM SIGMOBILE Mobile Computing and Communications Review*, 11(3):59–70, 2007.

[32] J. Parker, D.S., G. Popek, G. Rudisin, A. Stoughton, B. Walker, E. Walton, J. Chow, D. Edwards, S. Kiser, and C. Kline. Detection of mutual inconsistency in distributed systems. *IEEE Transactions on Software Engineering*, SE-9(3):240 – 247, May 1983.

[33] S. Pavlopoulos, E. Kyriacou, A. Berler, S. Dembeyiotis, and D. Koutsouris. A novel emergency telemedicine system based on wireless communication technology-ambulance. *IEEE Transactions on Information Technology in Biomedicine*, 2(4):261 –267, 1998.

[34] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser. A parsimonious model of mobile partitioned networks with clustering. *Communication Systems and Networks and Workshops, 2009. COMSNETS 2009. First International*, pages 1–10, 2009.

[35] I. Rhee, M. Shin, S. Hong, K. Lee, S. J. Kim, and S. Chong. On the Levy-Walk Nature of Human Mobility. *IEEE/ACM Transactions on Networking*, 19(3):630–643, 2011.

[36] V. Srinivasan, M. Motani, and W. T. Ooi. Analysis and implications of student contact patterns derived from campus schedules. In *Proceedings of MobiCom*, 2006.

[37] D. Williams and Naval Health Research Center (U. S.). *Tactical Medical Coordination System : (TacMedCS)*. Naval Health Research Center, San Diego, CA, 2007.

[38] W. Zhao. A new mac layer routing protocol for infrastructure wireless mesh networks. Master's thesis, Technical University of Denmark, 2008.

[39] X. Zhao, A. Rafiq, R. Hummel, D.-Y. Fei, and R. C. Merrell. Integration of information technology, wireless networks, and personal digital assistants for triage and casualty. *Journal of Telemedicine and E-Health*, 12(4):466–474, 2006.