Patrick Lieser, Flor Alvarez, Paul Gardner-Stephen, Matthias Hollick, Doreen Bhnstedt. *Architecture for Responsive Emergency Communications Networks*. In Proc. 7th Global Humanitarian Technology Conference (GHTC), IEEE, 2017. ISBN:978-1-5090-6046-7

Architecture for Responsive Emergency Communications Networks

Patrick Lieser[†§], Flor Alvarez[*§], Paul Gardner-Stephen[‡], Matthias Hollick[*] and Doreen Boehnstedt[†]
[†]Multimedia Communications Lab, TU Darmstadt, Germany
Email: {patrick.lieser, doreen.boehnstedt}@kom.tu-darmstadt.de
[*]Secure Mobile Networking Lab, TU Darmstadt, Germany
Email: {falvarez, mhollick}@seemoo.tu-darmstadt.de
[‡]College of Science & Engineering, Flinders University, Australia
Email: paul.gardner-stephen@flinders.edu.au
([§]Co-first authors: These authors contributed equally to this work)

***Abstract*—Self-organizing Mobile Ad-hoc Networks (MANETs) based on Delay Tolerant Networking (DTN), are powerful tools for maintaining or reestablishing telecommunications following disasters and other infrastructure disrupting events. However, such networks typically have very limited bandwidth compared with infrastructure-based networks, with the practical effect that they cannot satisfy every demand placed upon them. Thus, if the most critical traffic is to be delivered, and in a timely manner, some form of filtering or prioritization is needed. This paper sets out an architecture for solving this problem, and presents supporting simulation and field results. The architecture is built using the input of several emergency and disaster response organizations, to ensure that the key services required by citizens post-disaster were incorporated. Reflecting the dynamic nature of post-disaster communications needs, as identified in the survey, the architecture provides a framework in which arbitrary prioritization policies can be defined, and redefined, so that the humanitarian utility of a network can be maximized according to the prevailing situation and requirements. A proof-of-concept implementation is presented, yielding orders of magnitude reduction in message delivery latency in both simulation and in a field trial of an existing disaster communications system.**

I. Introduction

The many recent disasters, such as bush fires [1], earthquakes [2], and floods [3], have demonstrated the difficulty of organizing affected populations, especially when such disasters occur in urban environments. A common characteristic of these disasters is the loss of mobile telecommunications capability [4], which hinders citizens and responders alike, highlighting the need for resilient communications solution.

Typically, in such situations, the most critical personal communication needs focus on the exchange of small but vital data, such as SOS messages, telling family and friends that you are safe, or sharing situational awareness [5]. DTN-MANET, a combination of Delay Tolerant Network (DTN) and Mobile Ad-hoc Network (MANET), is well-suited to enable such communications post-disaster: it does not depend on any fixed or conventional infrastructure, or end-to-end connectivity, but instead it can be rolled out as needed and is easily adaptable or relocatable [6], [7]. Thus citizens and responders can build spontaneous communication networks using their mobile devices and communicate directly, as depicted in Figure 1.

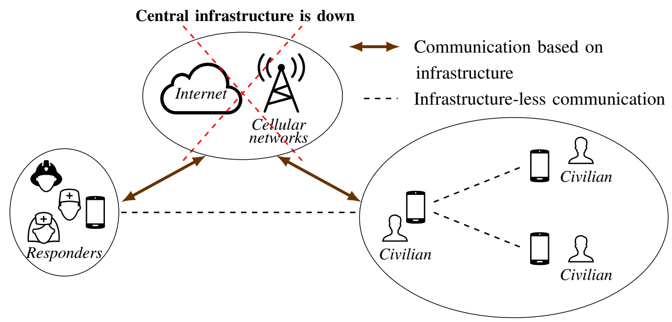

Fig. 1. Communication between users in a self-organized post-disaster response network.

However, delay tolerant and mobile ad-hoc networks provide only the communications mechanism, but not the data services built on them, e.g., SOS messages in disaster situations. Also, DTN-MANETs are likely to offer only a fraction of the bandwidth of infrastructure-based networks. This relative scarcity of bandwidth is exacerbated by the surge in demand for telecommunications typically following a disaster [8]. Thus, while these post-disaster communication systems may not fully replace cellular communication infrastructure, they can provide an additional or backup communication channel that is cheaper, more capable and more ubiquitous than two-way radio – provided that bandwidth use can be optimized sufficiently to allow delivery of at least critical communications.

In delay tolerant and mobile ad-hoc networks, each node generates and manages its information and the information received from the neighbor nodes locally. For disaster response, it is necessary to design DTN-MANETs, where each node locally assesses data priority, so that it can expedite high priority data (Figure 2). The goal is that every node, at every opportunity for transmission, selects the highest-priority item to transmit. The challenge is then reduced to two key components: (1) devising optimal data prioritization policies – which may be different for different disasters – that allows for the relative priority of data items to be computed, and; (2) the creation of an architecture that facilitates such flexible definition, modification and execution of communications priority policies. The focus of this paper is on the the second component, as it forms the foundation for enabling such optimal infrastructure independent post-disaster networks.

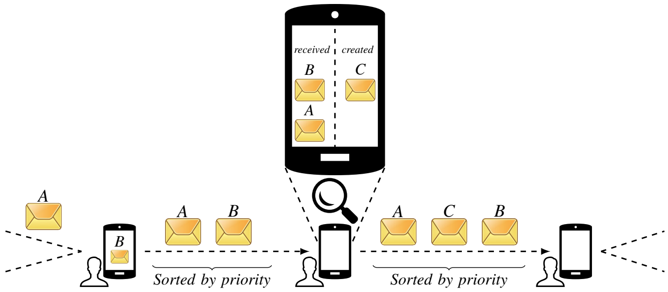

Fig. 2. Basic example of a disaster communication using data prioritization in delay tolerant and mobile ad-hoc networks (DTN-MANETs).

II. Contributions

This paper provides the following contributions:

1. A summary of relevant related work in Section III.
2. A representative list of disasters in Section IV, highlighting common factors.
3. Enumerating and analyzing the information services vital to disaster response in Section V, based on a survey of German fire and disaster response organizations.
4. The formulation of a system architecture based on these information services in Section VI. The architecture has been specifically designed for flexibility, to allow the implementation of continuously evolving communications priority policies.
5. Finally, a demonstration of the benefit of our architecture, by implementing it in the Serval Mesh open-source disaster communications system [9]–[11] in both simulations and a field trial. The results in Section VII show the clear benefit of such a prioritization scheme, and the effectiveness of even simple prioritization rules. Specifically, the results demonstrate orders of magnitude reduction of SMS-like message delivery latency in a heavily loaded post-disaster communications system based on delay tolerant and mobile ad-hoc networks.

III. Related Work

In this section, we investigate and summarize relevant post-disaster communication systems based on DTN-MANETs. We focus on the data exchanged between the devices as well as how these systems use prioritization schemes to deal with the limitations imposed by delay tolerant and mobile ad-hoc communications.

1) Twimight [12]: is a twitter-based android application that allows communication between twitter users via Bluetooth, using an epidemic routing protocol. The messages are stored locally and are sent to the Twitter servers whenever connected to the internet. The system uses fixed prioritization scheme, based on type of message. One disadvantage, is that users need a pre-existing Twitter account, or an internet connection to create one. Additionally, Twimight supports only hashtags, i.e., it lacks user and group communications.

2) SOS-Cast [13]: is an Android application that helps responders to find trapped persons. A user broadcasts an SOS message, including their name, physical condition and location. The priority of each message is defined by a static priority policy, incorporating factors such as sender identity, location and signal strength indication. The messages are then relayed via civilian phones until they reach a responder device. Communication is only supported from civilians to responders.

3) TRIAGE [14]: is a framework for emergency response communications that calculates message priority by considering the user context, the role of the sender, and the content of the message. This information is also used to schedule and to prioritize critical messages at the TCP level. However, Triage requires a connection with a central infrastructure to determine the extent of the network congestion and its main focus is communication between responders.

4) Serval Mesh [9]–[11]: is a general purpose and infrastructure independent DTN-MANET communications system. It allows both one-to-one, and one-to-many communications, which includes support for automatic low-speed long-range packet radio communications. It has, however, historically lacked a flexible means of defining message priority.

A. Discussion

Although there are several projects focused on the design and implementation of infrastructure independent post-disaster communication systems, all lack a flexible prioritization architecture. A general or reference architecture which can accommodate multiple scenarios or requirements and prioritize dynamically is missing. Thus, rather than proposing another specific system, we instead present a multi-layered reference architecture, that can be incorporated into existing projects, and which is described in detail in Section VI.

IV. Use Cases

In this section, we summarize relevant disaster scenarios from the last few years to identify the most common communications needs from the perspective of both organizations and individuals. Additionally, we highlight the main communication issues that occurred in these scenarios.

1) Earthquake & Tsunami: In April 2015, Nepal suffered a magnitude 7.8 earthquake, causing significant damage to the local telecommunications infrastructure. The disruption of communications complicated relief efforts. It hindered the coordination of the help effort, slowing the response, especially during the crucial first hours [15]. Similarly, following the magnitude 9.3 Sumatran earthquake of 2004 [16], warning and relief efforts were impaired by damage to infrastructure and lack of communications [17]. Particularly following the Nepal earthquake, citizens played an important role in reducing these effects, assisting relief efforts through collecting, disseminating and exchanging information and news about the ongoing situation in the disaster area via social networks. Social media was also used to search for missing people or relatives, and to reassure others of their own safety.

2) Hurricane Katrina: In August 2005, hurricane Katrina, one of the five deadliest hurricanes in the history of the United States, caused huge economic damage and resulted in the isolation or death of thousands of people. One of the main contributors was the lack of functional communications [18]. Only a very limited number of communication channels were available to inform the affected population about the urgency to evacuate. The pervasive and widespread communications failures also substantially hampered relief efforts. For example, responders were forced to door-knock to inform residents that they needed to evacuate. Where communications were available, civilians used social media to share information about the situation in the affected areas, as well as to inform others of their most urgent needs following the disaster.

3) Black Saturday Bush Fires: In 2009, the Black Saturday bush fire affected a widespread area in the southern Australian state of Victoria [1]. This was Australia's worst bush fire since the Ash Wednesday fires of 1983, resulting in 173 fatalities, with communications services unavailable in many areas [19]. The lack of information about the accessibility of the affected areas greatly delayed the relief efforts. Before they could act, responders had to collect information about the impact of the disaster to facilitate/enable access [20]. Furthermore, affected regions were unable to receive warnings and evacuation instructions in a timely and reliable manner, thus compounding the situation. Nonetheless, the collaboration of local communities, local and international organizations helped to provide support to people affected by the disaster. For instance, the Red Cross registered the names of affected people to collect information about their safety and enabling inquiries about missing persons [21].

4) War and Unrest: War and civil unrest also often disrupt communications infrastructure, through either damage or other actions of the belligerents. Such unrest also acts to impair the development and extension of telecommunications infrastructure. South Sudan is an example of this, where years of civil conflict and warfare have acted to prevent investment in telecommunications infrastructure [22].

A. Summary

While the use-cases are varied, they tend to exhibit a number of common factors, such as loss of communications capacity, or the isolation of people and communities from one another. Indeed, if we consider these and several other representative factors for the above use-cases, we find that they almost all apply to every use-case (Table I), although differences may arise the in the relative significance of each factor.

TABLE I. OVERVIEW OF RELEVANT DISASTER SCENARIOS AND THEIR ISSUES.

	Nepal / Sumatra Earthquake / Tsunami	*Hurricane Katrina*	*Black Saturday Bush Fires*	*War and Unrest*
Loss of Communications	●	●	●	◑
Isolation of People	●	●	●	◑
Response Difficulties	●	●	●	●
Use of Social Media	●	●	◑	●
Collaboration of Citizens	●	●	●	◑
Search for Missing People	◑	●	◑	◑
Lack of Information	●	●	●	◑

● fully applies, ◑ partially applies, ○ does not apply

V. RELEVANT SERVICES

The aforementioned use-cases are only examples of how users and their needs influence the emergency response networks dedicated to helping them. Self-organized public networks may be used in varying scenarios by people with very different requirements, and types of information to be exchanged. Before defining the relevant services, it is thus important to consider the role of people in disasters, as well as to identify possible forms of data dissemination.

A. Communication pathways

The roles of people during a disaster are complex. Many beneficiaries of help are at the same time responders. A person may, for example, be a beneficiary of food and water, but then assist the response efforts by searching for survivors or restoring telecommunications services. Such persons may provide help in both individual and institutional capacities. Furthermore, the institutions or organizations which they serve may either be pre-existing ones, e.g., a cellular carrier or national disaster management organization, or a newly formed ad-hoc relief organization [23]. Thus, it is extremely difficult to cleanly divide the roles of people in a disaster zone.

Together with our evaluation of the cases surveyed, this led us to the understanding that communications between civilians in a disaster zone is both of vital importance, and often particularly vulnerable, as civilians typically do not possess their own dedicated communications capacity, in contrast to many established relief organizations.

Therefore our approach concentrates on communications tools that ordinary civilians can make use of in the wake of disasters, as depicted in Figure 3. As such, our focus is the facilitation of civilian-to-civilian (C2C) communications. Organizational or institutional communications are also explored, but only in so far as they can be directed at, or received by civilians, i.e., organization-to-civilian (O2C) or civilian-to-organization (C2O) communications. Communications between organizations (O2O) is outside of the scope of this paper.

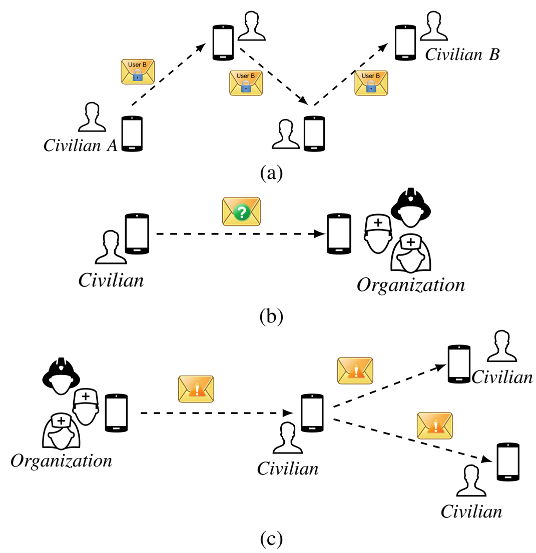

Fig. 3. The three communication pathways: (a) civilian-to-civilian (C2C), e.g., unicast private message, (b) civilian-to-organization (C2O), e.g., anycast help request, (c) organization-to-civilian (O2C), e.g., broadcast warning messages.

For these communication pathways, there are several forms of dissemination available: *unicast* for a particular recipient; *broadcast* for everyone on the network; *multicast* for a specific group in the network; *anycast* for at least one recipient in a group; and *geocast* for a group within a certain vicinity.

B. Services

Based on the use-cases summarized above, and the input of several responder organizations such as German Fire Departments, and the German Federal Office of Civil Protection and Disaster Assistance, we define the following list of relevant services:

1) SOS Emergency Messages: This service allows people to send an urgent request for help to responders, as fast as possible (C2O). The message may additionally be sent to neighboring nodes (C2C) that can act as first responders. This service works as an addition, and not as a replacement, to national emergency numbers, such as 112 in Europe, 000 in Australia and 911 in the USA.

2) I am Alive Notifications: I am Alive Notifications enable the affected population to report their status information, e.g., location, health, and felt needs to other users in the post-disaster network (C2C). As mentioned in the previous section, social networks like Facebook, and NGOs like the Red Cross, implement such services on demand, mostly based on websites hosted on central servers. These services are rarely integrated, resulting in fragmentation of information, and often requiring users to submit their information on multiple systems, further straining communications infrastructure.

3) Person-Finder: Person-Finder provides the counterpart to I am Alive Notifications: the possibility to ask about people assumed to reside in the area around the incident (C2C). To facilitate this, the service should allow searching for people based on different information, e.g., last known location or via photograph. Geographic forwarding schemes may be employed to limit and refine searches to the assumed location of an individual.

4) Situation Assessment: This service allows the affected population to report observations from the disaster area, such as damage reports or availability of supplies, to either responders or the affected population (C2C & C2O).

5) Information/News: Information/News services allow responders to make announcements regarding currently existing, or potentially evolving hazards in a specific area to the public (O2C). This service could use various dissemination modes, targeting groups, individuals, areas, or any combination.

6) Resource Market Registry: This service is used to match requests for resource, e.g., requests for fuel, energy, water or medical supplies, to respective offers from the affected population. This service provides a tool for self-organized resource sharing based on information about needs and requests. The information should be exchanged among the affected population (C2C), but only in specific regions, in order to prevent unnecessary information transmission, and thus minimize the required communications capacity.

7) Tasking: The Tasking service is similar to the Resource Market Registry Service, but focuses on human resources, i.e., enabling responders or the affected population to recruit and manage personnel in achieving particular relief initiatives of individuals (C2C or organization O2C).

8) Messaging Services: Messaging service allows private messaging similar to SMS between two parties, enabling the affected population to communicate with family, friends, or others for any necessary purpose (C2C).

C. Summary

Most of these services are individually provided by existing commercial solutions. However, as of today, all of them are highly dependent on centralized infrastructure, i.e., are based on a client-server architecture. For example, there are websites or apps that supply interactive maps for actual or potential disasters, such as hurricanes or tsunamis (Disaster Alert [24]). Also, many solutions enable users to report an incident and to get feedback about the current status of service restoration (FEMA App [25]). However, although they exist for institutions and organizations, we are not aware of citizen-oriented Resource Market Registries or Tasking Services. Even if they were available, dependence on communications infrastructure would remain an obstacle to their use.

VI. System Architecture

In this section, we describe a general architecture for post-disaster communication systems. We focus primarily on post-disaster networks based on delay tolerant and mobile ad-hoc networks (DTN-MANETs) where no fixed infrastructure is available.

In such networks, each node can act as data source, destination or relay station, and thus needs to decide whether data is forwarded, stored or discarded. Our architecture is based on ubiquitous mobile devices equipped with a variety of common sensors. The sensory capabilities of such devices can provide helpful information about the extent or severity of damage at the site of the disaster [26] as well as the status of the device's owner through activity recognition [27].

Each device, however, has constrained resources such as battery capacity and communications modes and bandwidth. Additionally, in DTN-MANETs, where the interconnection time between nodes is not predictable, an end-to-end communication channel cannot be guaranteed.

Therefore, this kind of system requires appropriate mechanisms that facilitate optimal data exchange under any circumstances. Our architecture addresses these challenges, including the following key points: the services described in the previous section, the information gained from the device's sensors, and prioritization mechanisms to allocate different levels of data priority.

The architecture is not limited to a specific communication technology, and can support existing network protocols and physical interfaces required for direct device-to-device communication.

As depicted in Figure 4, our architecture is structured in three main layers: the Service-, Intelligence- and Communication layer.

A. Service Layer

The Service Layer generates the information utilized for the attached Intelligence Layer. This layer consists of a Sensor- and a Message/Data Module.

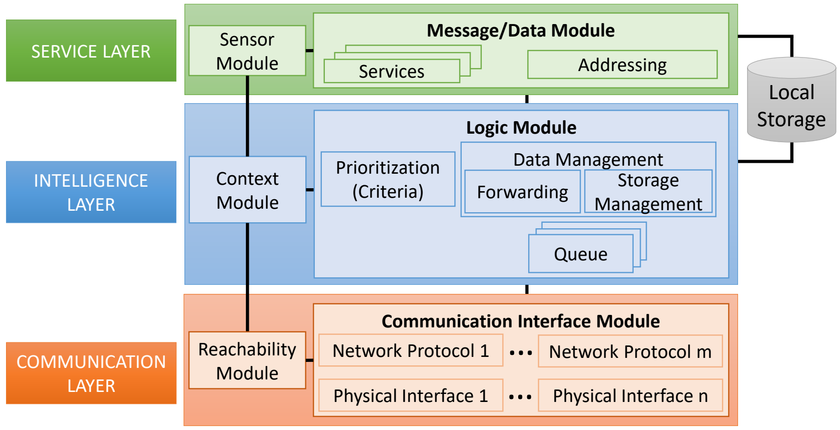

Fig. 4. Overview of the proposed architecture for post-disaster communication systems.

1) Message/Data Module: This module fulfills two main tasks. First, it generates a new message of a given service class, either from the list in Section V, or corresponding to a novel service class. Messages can be either user supplied, or automatically generated, for example, in response to sensor readings.

Second, this module sets the destination of each message based on the selected service, and using one of the dissemination modes outlined above.

2) Sensor Module: This module collects and manages the sensor information relevant to any of the services within the architecture. This information is used by the Message/Data Module to add specific information to a new message, e.g., adding the user's location to a resource market service request or to a private text message to a family member.

B. Intelligence Layer

The Intelligence Layer represents the core of our architecture. This layer consists of the Logic and Context modules. The Logic Module is responsible for the data management and the implementation of prioritization mechanisms. The Context Module generates the context information used by the prioritization mechanisms.

1) Context Module: This module collects information provided by the Sensor and the Reachability Modules. The Reachability Module provides information about the network conditions, such as which devices are reachable, network link quality, or the recent network activity of this or other users. All collected information is passed to the Logic Module, so that it can be acted on.

2) Logic Module: The functionality of this module is split into three main components: Queue, Data Management and Prioritization.

The ***Queue*** collects the data to be sent to the Communication Layer, and uses the priorities computed by the Logic Module to specify the order of dispatch to the Communications Layer. The queue may be a single queue, or may be a collection of multiple queues, depending on the complexity of the priority scheme selected by the Prioritization Component.

The ***Data Management*** Component is responsible for forwarding, storing or discarding messages. Messages that it accepts from the Service Layer for forwarding are passed to the Queue. However, it first verifies the validity of a message, e.g., examining any time-to-live parameter or cryptographic signature. For messages received from the Communication Layer, it also checks whether the message should be forwarded to the Service Layer, because it is of local relevance, or whether it should only be forwarded to other nodes, or whether it should be ignored, because it is for some reason irrelevant, for example, outside of a self-described geographical or temporal bounding box.

Prioritization consists of a mechanism that can assign different priorities to the data, either statically, adaptively, or in a combined mode, depending on the available information and the desired sophistication of prioritization. Prioritization of data plays an important role in delay tolerant and mobile ad-hoc post-disaster communication systems, where each node has limited resources and is responsible for the processing of the information generated on the entire network.

In static mode, messages can be prioritized by any defined metric, such as the size of a message, number or identity of recipients, or the message or service type.

While the proposal of specific metrics and rules is beyond the scope of this paper, it is apparent that it is possible to attempt to optimize against arbitrary utility metrics, for example, maximizing the number of messages delivered per unit time, minimizing the latency for some class of messages, for example SOS messages, or maximizing the dissemination of critical information. It may be possible to compose rules that improve more than one of these utility metrics. Where such rules are possible, it makes sense to pursue them. One possible example is to prioritize shorter messages ahead of longer messages, which would maximize the number of messages delivered, while simultaneously minimize the latency for SOS messages, provided that such messages are among the shortest, which seems a reasonable assumption.

For adaptive prioritization, context information, such as reachability of message recipients, user activity, battery level, position or network conditions, can be used to decide the relative priority of messages. We describe two such examples.

First, if a node has low battery level, it may stimulate prioritization of messages addressed to it by others, so that the messages can, hopefully, be received by that node before its battery is depleted. Similarly, message traffic to nodes that are reachable on the network might be given a higher priority, since the probability of delivery of such messages is likely to be higher. Additionally the GPS-Location can be used as context information to prioritize messages which are generated in or are addressed to a certain area, e.g., the epicenter of an earthquake.

A second scheme to adaptively prioritize messages would be to monitor network traffic, to discover which message types are currently trending. This is relevant because the importance of message types can vary greatly during the course of a disaster. Reports from past disasters, such as the Yushu Earthquake 2010 in China [28], confirm the varying composition of message traffic during a disaster. For example, immediately following the earthquake, situation reports were prominent. Later, once people understood the situation, help requests and messages coordinating relief efforts came to the fore.

Whatever the means of adapting message priority, the result is to improve the agility of the resulting system, so as to maximize the utility of the constrained communications resources, in the face of the dynamic nature of disasters.

All messages, whether they were just created by the Service Layer or forwarded from another node, will be sent to the Queue Component.

The particular ***Prioritization Criteria*** for messages can be any arbitrary function that the operator of a device wishes. In the experiments in this paper, a simple size-based criteria was used. That is, smaller messages were given higher priority over larger messages. In a conventional routed network, packets may be dropped when congestion occurs. In contrast, in a prioritized transmission regime, low-priority communications are simply delayed, possibly for an indefinite period, before delivery, and only if storage on a node is exhausted, will they be dropped to make space for higher-priority communications. That is, prioritization is at worst equivalent to conventional routing, and at best, provides for late delivery where conventional routing would result in non-delivery.

C. Communication Layer

The Communication Layer coordinates the ad-hoc communications though management of device-to-device communication with neighboring nodes, and consists of two modules: First, the Communication Interface, which selects and couples the routing protocols and physical interfaces necessary to communicate between devices. Second, the Reachability Module provides information about the network to the Context Module.

1) Communication Interface Module: The Communication Interface Module consists of two main components: the Network Protocol and Physical Interfaces.

The ***Network Protocol Interface*** deals with the logic of the communication itself. It is responsible for tasks such as deciding how much information will be exchanged between two devices during an encounter. This component chooses a suitable routing protocol to be used for any communication. Its selection is influenced by factors that include the current network situation.

In a disrupted scenario where the nodes are only temporarily available, solutions such as Serval Rhizome [9], Prophet [29], or Epidemic [30], may be the desired approach.

For very low bandwidth links, such as packet radio, it may make sense to instead have the radio layer request data packets whenever it is able to transmit a packet, so that the prioritization decision is always made without the elapse of long periods of time before transmission.

In a more stable scenario, where the nodes are frequently available and end-to-end communication is possible, an ad-hoc routing solution, such as B.A.T.M.A.N [31], AODV [32], or Serval MDP [11], could enable more efficient data transfer.

The ***Physical Interface*** selects the most suitable interface for the respective physical data transmission. The architecture can be used with any device-to-device-capable physical interfaces available on the node, e.g., Bluetooth or Wifi-Direct.

2) Reachability Module: This module collects information about the current network status, e.g., available bandwidth or the number of neighboring nodes. This can be used by the Intelligence Layer, to allow the Logic Module to prioritize transmissions based on these factors.

D. Architectural Considerations

There are two concepts that require special consideration when designing post-disaster communication systems based on our architecture: network neutrality and security.

1) Network Neutrality: For normal non-disaster network conditions, network neutrality is highly desirable to avoid a variety of discriminatory and anti-competitive behaviors. That is, networks should not prioritize any data or communications over any other [33]: However, in a disaster situation, the scarcity of bandwidth and compelling humanitarian needs may necessitate rationing network resources, similar to other resources, such as food, water and shelter. That is, network equitability or fairness may be more important than neutrality.

2) Security: Security does not lose importance during disasters. Therefore, care should be taken when implementing this architecture, so as to maximize the security properties of the resulting networks. This may take many forms, such as verifying the authenticity or integrity of messages [34], [35]. Of particular concern, is ensuring the prioritization rules are not subject to spoofing or other forms of abuse. For example, if a prioritization rule allowed an SOS flag to be arbitrarily attached to messages to increase their priority, it is likely that users may mark non-SOS communications as SOS, in order to expedite delivery. The creation of prioritization rules that are robust against such abuses are outside the scope of this paper.

VII. Proof of Concept

To test the feasibility of our architecture, the proposed prioritization architecture was implemented in the Serval Mesh open-source diaster communications system [11]. We extended the Serval Mesh to include a flexible content prioritization system for the Serval Rhizome [36], [37] delay-tolerant networking system.

This implementation consisted of calculating and applying a dimensionless priority value to every data bundle, based on: (1) meta-data, e.g., the type and size of the content, who sent it or to whom it is addressed; and (2) the peers a device is in contact in, and which bundles they have. This allows for the most important prioritization, i.e., the suppression of transmission of already received content.

A. Serval Rhizome, LBARD and Serval Mesh Extenders

Serval Rhizome is a store-and-forward based communications protocol designed for disaster communications. The basic data unit in Rhizome is a Rhizome Bundle, which consists of a (possibly empty) file, and an associated manifest. The manifest is simply a <1KB file containing the bundle's meta-data, such as sender, recipient, service class, size and version. This meta-data can be used to calculate the priority of a bundle. Serval Rhizome is typically implemented by a node maintaining a database of Rhizome Bundles, i.e., a Rhizome Database.

Low-Bandwidth Asynchronous Rhizome Delivery (LBARD) is a variant of Serval Rhizome designed specifically for very low data rate links, i.e., links of <1KB / second. LBARD simultaneously operates two stages of communications: First, the Tree-Sync protocol synchronizes bundle inventories, i.e., the lists of Rhizome Bundles held by two communicating nodes. Second, the exchange of Rhizome Bundles. The operation of LBARD is explained in more detail in a companion paper [38]. Serval Mesh Extenders are low-cost robust outdoor disaster communications relay devices, that include Wi-Fi, an RFD900+ UHF packet radio [39], a Serval Rhizome database and an LBARD instance. These combined capabilities allow smart-phones in the vicinity of a Mesh Extender to communicate over thousands of meters with other smart-phones via the low-bandwidth Ultra-High-Frequency (UHF) packet radio. The low effective bandwidth of the UHF radio link, typically <1KB / second, necessitates effective prioritization of Rhizome Bundles in order to deliver acceptable performance. The Serval Mesh Extender is described in more detail in a companion paper [40].

All elements of the Serval Mesh are freely available under open-source licenses.

B. Experimental Design

The Serval Low-Bandwidth Asynchronous Rhizome Delivery (LBARD) program was used on prototype Serval Mesh Extender devices. Post-mortem examination of the Mesh Extenders revealed that each was carrying more than 1,300 Rhizome Bundles of varying sizes (1,287 0kB to 1kB, 24 1 to 10kB, 7 10kB to 100kB, 1 100kB to 1MB, 7 1MB to 5.8MB, total size >58MB). Each Mesh Extender was equipped with LBARD connected to an RFD900+ UHF packet radio operating at 921MHz, and providing approximately 1kB/second throughput.

Another version of LBARD was crafted incorporating a prioritization function, with the following rules: (1) Content already delivered had the lowest priority, and (2) Smaller Rhizome Bundles were prioritized over larger ones. The goal was to prioritize delivery of new small messages.

1) Experiment 1: Simulation: A set of experiments was performed, using the simulation framework that forms part of the Serval Mesh test framework. This framework simulates the UHF packet radio communications between Mesh Extenders. The Rhizome databases were populated with between one and 2,048 Rhizome bundles. Each experiment saw a single short message injected into the Rhizome database of one node, and the delivery time of that message was measured for each experiment, both with and without prioritization enabled.

2) Experiment 2: Field Test: In addition to the simulation work, field experimentation was carried out in the Arkaroola Wilderness Sanctuary in a remote location in Outback Australia, chosen for being representative of disaster zones. The Mesh Extenders were operated, connected via Wi-Fi to Android smart-phones running the Serval Mesh App version 0.93 for Android. These experiments consisted of two parts.

First, the Mesh Extenders were located in close proximity, allowing the packet radios to operate with negligible packet loss, and LBARD was operated without the prioritization function. The expectation was that no messages would be delivered within a reasonable period of time, due to the lack of prioritization.

Second, the Mesh Extenders were operated with prioritization enabled, located approximately 4km apart at the Sir Mark Oliphant Observatory (30°18'29.59"S 139°20'16.20"E) and Coulthard's Lookout (30°16'22.79"S 139°20'25.78"E) in Arkaroola as depicted in Figure 5.

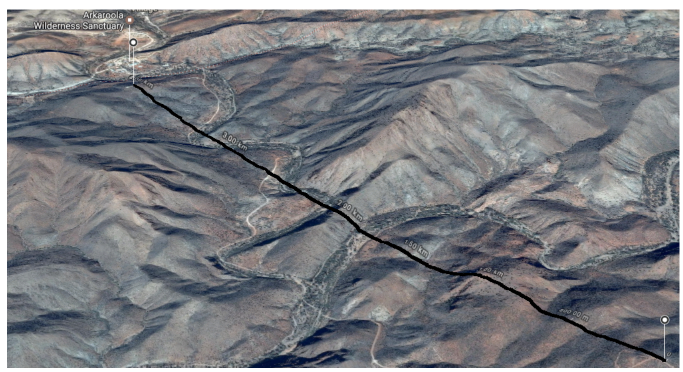

Fig. 5. Perspective view of the 4km communications path for the experiments. Rendered using Google Maps.

Figure 6 shows the vehicle-mounted Mesh Extender prototype that was used at the Lookout, together with the view towards the Observatory (marked by the arrow). The purpose of the experiments was to measure message delivery latency over a packet-radio hop in a real-world environment, including delivery confirmation, i.e., the sequential transmission and reception of one Rhizome bundle in each direction.

C. Results

1) Simulation: For the simulation experiments, Figure 7 presents the aggregated results for message delivery times, with varying numbers of bundles present in the Rhizome database of the sending node. For very small numbers of bundles, there is little difference in delivery time. This is expected, as there can be at most only a few KB of data in a few bundles that

Fig. 6. Coulthard's Lookout link end, looking towards the Sir Mark Oliphant Observatory link end in Arkaroola, which is indicated by the arrow (4km Line-of-sight distance).

can possibly be scheduled before the bundle to be transmitted. Then, as the number of bundles increases, the non-prioritized delivery time grows very rapidly, in an approximately linear relation to the number of bundles. The variation in delivery time reflects the random order in which the bundles are queued for transmission, in the absence of any prioritization. In contrast, when prioritization is enabled, delivery latency grows sub-linearly with the number of bundles, as the prioritization ensures that the injected bundle is always delivered first. The growth in delivery time reflects the time spent by LBARD to determine which bundles each node already possesses.

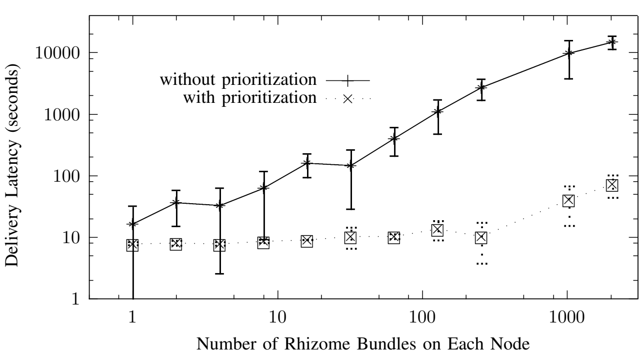

Fig. 7. Delivery time (including standard deviation) for a short message bundle, with and without prioritization ($n = 4$).

2) Field Test: For the real-world experiments in Arkaroola, for those experiments where prioritization was enabled in LBARD, message delivery, including confirmation of delivery notification, occurred in <1 minute. In contrast, without prioritization, no messages were delivered within 1 hour.

D. Discussion

The experimental and simulation results are in agreement: Prioritization allows for delivery of a high-priority message to occur within tens of seconds, even when 1,000s of bundles are in transit, and when using a slow UHF packet radio. In contrast, when there is no prioritization of traffic, the delivery latency grows approximately linearly with the number of bundles in transit, rapidly rendering the system unresponsive.

VIII. Conclusion

In this paper we have set out the case for a flexible architecture for disaster communications, that can support a wide range of use-cases and message services, and that can facilitate the flexible prioritization of message traffic within a disaster communications network, to accommodate the dynamically evolving demands that are placed on such networks, reflecting the complex reality of disaster situations.

Based on these insights, we proposed a general system architecture for post-disaster communication systems, considering both the needs of users, as well as the strengths and weaknesses of spontaneously established delay tolerant and mobile ad-hoc networks. This architecture has been designed to be agile and flexible enough to implement various and changing communications priority policies, consistent with the identified need. This architecture is not limited to a specific communication technology, and can support existing and future network protocols as well as new device-to-device physical communications interfaces.

We evaluated the practicability of our architecture, by incorporating it into the open Serval Mesh open-source disaster communications system. The results of simulations and field experiments confirmed that the addition of flexible prioritization, based on our proposed architecture, can deliver orders of magnitude reduction in latency in the face of congested channels, that are the reality in many disaster situations. Indeed, the implementation of such a priority scheme made the difference between the Serval Mesh being usable or not in the field conditions under which it was tested, suggesting that the proposed architecture has the potential to make a strong contribution in improving communications post-disaster.

Acknowledgment

This work has been funded by the German Federal Ministry for Education and Research (BMBF, Germany) within the SMARTER project and the Software Campus project "DisasterSense" (Support Code: 13N13406 & 01IS12054) and by the LOEWE initiative (Hessen, Germany) within the NICER project. The authors would like to acknowledge the support of USAID (USA), the NLnet Foundation (Netherlands), Radio Free Asia (USA), the Shuttleworth Foundation (South Africa), the Arkaroola Wilderness Sanctuary (Australia), and the Australian Department of Foreign Affairs and Trade's Foreign Aid Program (Australia).

References

[1] J. Whittaker, K. Haynes *et al.*, "Community safety during the 2009 australian'black saturday'bushfires: an analysis of household preparedness and response," *International Journal of Wildland Fire*, vol. 22, pp. 841–849, 2013.

[2] K. Goda, T. Kiyota *et al.*, "The 2015 gorkha nepal earthquake: insights from earthquake damage survey," *Frontiers in Built Environment*, vol. 1, p. 8, 2015.

[3] M. Nohara, "Impact of the great east japan earthquake and tsunami on health, medical care and public health systems in iwate prefecture, japan, 2011," *Western Pac Surveill Response J*, vol. 2, 2011.

[4] L. Palen and S. B. Liu, "Citizen communications in crisis: anticipating a future of ict-supported public participation," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2007, pp. 727–736.

[5] M. Zook, M. Graham *et al.*, "Volunteered geographic information and crowdsourcing disaster relief: a case study of the haitian earthquake," *World Medical & Health Policy*, vol. 2, no. 2, pp. 7–33, 2010.

[6] I. F. Akyildiz, W. Su *et al.*, "Wireless sensor networks: a survey," *Computer networks*, vol. 38, no. 4, pp. 393–422, 2002.

[7] S. K. Sarkar, T. G. Basavaraju, and other, *Ad Hoc Mobile Wireless Networks: Principles, Protocols, and Applications*, 2nd ed. CRC Press, Inc., 2013.

[8] F. Ranghieri, "Learning from megadisasters: lessons learnt from the great east japan earthquake and tsunami," *Economics and Policy of Energy and the Environment*, 2014.

[9] P. Gardner-Stephen, "The serval project: Practical wireless ad-hoc mobile telecommunications," *Flinders University, Adelaide, South Australia, Tech. Rep*, 2011.

[10] P. Gardner-Stephen, A. Bettison *et al.*, "The rational behind the serval network layer for resilient communications," *Journal of Computer Science*, vol. 9, no. 12, p. 1680, 2013.

[11] P. Gardner-Stephen, R. Challans *et al.*, "The serval mesh: A platform for resilient communications in disaster & crisis," in *Global Humanitarian Technology Conference (GHTC), 2013 IEEE*. IEEE, 2013, pp. 162–166.

[12] T. Hossmann, F. Legendre *et al.*, "Twitter in disaster mode: Opportunistic communication and distribution of sensor data in emergencies," in *Proceedings of the 3rd Extreme Conference on Communication: The Amazon Expedition*. ACM, 2011, p. 1.

[13] N. Suzuki, J. Zamora *et al.*, "SOSCast: Location Estimation of Immobilized Persons through SOS Message Propagation," in *Intelligent Networking and Collaborative Systems (INCoS), 2012 4th International Conference on*, sep 2012, pp. 428–435.

[14] F. Luqman, F. Sun *et al.*, "Prioritizing Data in Emergency Response Based on Context, Message Content and Role," in *Proceedings of the 1st International Conference on Wireless Technologies for Humanitarian Relief*, dec 2011, pp. 63–69.

[15] M. Wendelbo, F. La China *et al.*, "The crisis response to the nepal earthquake: Lessons learned," 2016.

[16] R. Paris and Wassmer, "Tsunamis as geomorphic crises: lessons from the december 26, 2004 tsunami in lhok nga, west banda aceh (sumatra, indonesia)," *Geomorphology*, vol. 104, pp. 59–72, 2009.

[17] T. Lay, H. Kanamori *et al.*, "The great sumatra-andaman earthquake of 26 december 2004," *Science*, vol. 308, pp. 1127–1133, 2005.

[18] US Senate, "Hurricane katrina: A nation still unprepared," *Rep. to the Committee on Homeland Security and Government Affairs, Washington, DC*, 2006.

[19] M. Bingemann, "Telstra says it's ready for the next black saturday," 2009. [Online]. Available: http://www.theaustralian.com.au/business/technology/telco-says-its-ready-for-the-next-black-saturday/news-story/f48e065f14fa8614f2a18165c9b279b8

[20] A. M. Gill and S. L. Stephens, "Scientific and social challenges for the management of fire-prone wildland–urban interfaces," *Environmental Research Letters*, vol. 4, 2009.

[21] Red Cross, "Issue 9 The Humanitarian Black Saturday Sticking Together," 2009. [Online]. Available: http://www.redcross.org.au/files/The_Humanitarian_9.pdf

[22] Emergency Telecommunications Cluster, "Current emergencies – south sudan," 2017. [Online]. Available: https://www.etcluster.org/countries/south-sudan

[23] D. J. Patterson, "Haitian resiliency: A case study in intermittent infrastructure," *First Monday*, vol. 20, no. 8, 2015.

[24] Pacific Disaster Center, "Disaster alert," 2017. [Online]. Available: http://www.pdc.org/solutions/tools

[25] Federal Emergency Management Agency, "Fema app," 2017. [Online]. Available: http://www.fema.gov/mobile-app

[26] H. Ma, D. Zhao *et al.*, "Opportunities in mobile crowd sensing," *IEEE Communications Magazine*, vol. 52, pp. 29–35, 2014.

[27] O. D. Lara, Labrador *et al.*, "A survey on human activity recognition using wearable sensors," *IEEE Communications Surveys & Tutorials*, vol. 15, pp. 1192–1209, 2013.

[28] Y. Qu, C. Huang *et al.*, "Microblogging after a major disaster in china: a case study of the 2010 yushu earthquake," in *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. ACM, 2011, pp. 25–34.

[29] A. Lindgren, A. Doria *et al.*, "Probabilistic routing in intermittently connected networks," *ACM SIGMOBILE mobile computing and communications review*, vol. 7, pp. 19–20, 2003.

[30] A. Vahdat, D. Becker *et al.*, "Epidemic routing for partially connected ad hoc networks," Tech. Rep., 2000.

[31] D. Johnson, N. Ntlatlapa *et al.*, "A simple pragmatic approach to mesh routing using batman," in *2nd IFIP International Symposium on Wireless Communications and Information Technology in Developing Countries, Pretoria, South Africa*, 2008.

[32] C. E. Perkins and E. M. Royer, "Ad-hoc on-demand distance vector routing," in *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1997, pp. 90–100.

[33] Federal Communications Commission, "The open internet," 2011. [Online]. Available: https://www.fcc.gov/general/open-internet

[34] M. Farb, Y.-H. Lin *et al.*, "Safeslinger: easy-to-use and secure public-key exchange," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 417–428.

[35] M. Schmittner and M. Hollick, "Xcastor: Secure and scalable group communication in ad hoc networks," in *2016 IEEE 17th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. IEEE, 2016, pp. 1–6.

[36] P. Gardner-Stephen, A. Bettison *et al.*, "Improving compression of short messages," *Int'l J. of Communications, Network and System Sciences*, vol. 2013, 2013.

[37] P. Gardner-Stephen, J. Lakeman *et al.*, "Meshms: Ad hoc data transfer within mesh network," 2012.

[38] G. Al-Nuaimi and P. Gardner-Stephen, "Efficient & effective use of low-bandwidth radio links for public & private disaster communications," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[39] RFDesign, "Rfd900 radio modem data sheet," 2013. [Online]. Available: http://files.rfdesign.com.au/Files/documents/RFD900\%20DataSheet.pdf

[40] P. Gardner-Stephen, "Productizing humanitarian telecommunications research: A case study of the serval mesh extender," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.