CNS: Content-oriented Notification Service for Managing Disasters

Jiachen Chen^{*†}, Mayutan Arumaithurai[‡], Xiaoming Fu[†], and K. K. Ramakrishnan[‡]

^{*}WINLAB, Rutgers University, NJ, U.S.A. jiachen@winlab.rutgers.edu

[†]Institute of Computer Science, University of Göttingen, Germany. {arumaithurai,fu}@cs.uni-goettingen.de

[‡]University of California, Riverside, CA, U.S.A. kk@cs.ucr.edu

ABSTRACT

Disaster management critically depends on timely and efficient communication. To better deal with an incident, authorities from different services (*e.g.*, fire, police) and jurisdictions need to work together in a new dynamically created team, different from their original organizational/administrative hierarchy. Unfortunately, existing solutions (*e.g.*, IP, or traditional telephony) are not well-suited to deal with such group communication due to the dynamic binding between *roles* and individuals, and mobility. A significant burden is placed on administrators to just establish and maintain necessary channels, distracting them from restoring order. To make things worse, since senders do not know which individual(s) to send to, information cannot reach the right people, delaying rescue efforts.

We propose CNS, leveraging the benefits of ICN to provide the essential communication for efficiently managing disasters. We first design a namespace enabling dynamic creation and evolution of incident related (sub-)namespaces to represent roles of first responders assigned to the disaster. This allows first responders to receive the appropriate information on a timely basis, with senders addressing the recipients based on their roles. Predefined namespace templates for disaster types minimize management overhead for establishing communication. We also find the need for a new enhanced forwarding rule to support such a recipient hierarchy.

We have developed a prototype demonstrating feasibility and efficiency. With the help of large-scale simulations and real-world disaster traces, we compare CNS with an IP-based solution. CNS can significantly reduce network load and latency in addition to the qualitative benefits of simplified operations, appropriate prioritization and security.

CCS Concepts

• **Networks** → **Application layer protocols; Naming and addressing;**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICN'16, September 26 - 28, 2016, Kyoto, Japan

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4467-8/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2984356.2984368>

Keywords

ICN; Disaster Management; Naming Structure; Notification

1. INTRODUCTION

We have witnessed an increasing number of disaster events, both natural events and ones caused by terrorists in the recent past. Most of these events, such as the 2005 London bombing (adversarial), the 2014 Kaohsiung gas explosion (accidental), and the 2015 Nepal earthquake (natural) point to a common trend in that these disasters often comprise multiple incidents occurring in different places around the same time. These events result in a massive need for first-response teams to manage the aftermath of the disaster, including rescue operations, dealing with emergencies and ensuring people are safe and appropriately informed to prevent panic. A detailed report of London Bombing [1] emphasized the need for enhanced communication capabilities and repeatedly recommended putting in place effective communications within *and between* the emergency services for such incidents. There is a clear need for cross-functional cooperation and collaboration across administrative and management boundaries to manage the disaster. Special teams of first-responders with different complementary expertise have to be dynamically formed, with a management and organizational structure that is different from their normal management and administrative hierarchy and responsibility. In fact such dynamically formed teams responding to the incident might also include members who are not part of the traditional first response service team (*e.g.*, the role of the underground control center in the London Bombing incident).

However, we observe that it is difficult to achieve the effective communication in the aftermath of a disaster with existing communication frameworks. The first difficulty lies in the fact that in order to send a message (either calling for help or providing information about disaster), the sender needs to know the specific individual(s) (and their phone numbers or eventually their IP addresses) who are dealing with the incident. As a result, it often causes messages to not go to the right people and results in confusion rather than deeper understanding of the disaster among the first responders. The consequence of such confusion is that the management of the event is inadequate, with commanders erroneously dispatching or not allocating resources appropriately. What makes matters worse is that, once the first-responders are mobilized, the lack of communication among those dynamically formed groups results in difficulty in correcting for errors (*e.g.*, redeployment) or sending them information in a timely, efficient manner as the

managers/commanders have to keep track of the position of each unit and send messages to the individuals separately. All these issues resulted in delayed response and poor outcomes for disaster management.

Timely information dissemination to the right recipient is important for disaster management. A communication framework based on IP – a location-dependent protocol – has inherent difficulties since the communication has to be based on the individual’s address. On the other hand, Information-Centric Networking (ICN) paradigms such as NDN [2], MobilityFirst [3] and XIA [4], among others, treat contents and names as first class entities. These solutions enable the access of information based on its name or identity, without regard to location. The ability to access and disseminate information with network support enables ICNs to deliver the desired information in a timely manner. Some ICN approaches, such as [5], provide enhancements to ICN to get efficiencies in delivering information using a “push” semantic (publish/subscribe and multicast). With such an approach, establishment of a group is convenient without having to first distribute group IP addresses before information is exchanged. These capabilities of ICN are highly desirable for communication in disaster management situations.

However, these approaches need to be further enhanced to satisfy the communication needs of disaster management, where context- and recipient-driven communication is needed. Senders need to address recipients based on their roles/persona (or context) rather than the individual(s) or their addresses. *E.g.*, a commander might want to send a command to all firemen dealing with the London Bombing at Aldgate site without the need to know each of them. Further, members should have a way to dynamically form specific teams (instantiate their roles) in order to send and receive messages efficiently. Of course, they need to retain their original administrative hierarchy for other purposes.

In this paper, we propose the introduction of two capabilities in ICN to achieve this functionality. The first is a flexible namespace that can represent the context (or organization) of the normal administrative hierarchy as well as the special recipient hierarchy needed for incident response. The namespace should be able to evolve (create, modify and revoke elements) dynamically during the lifetime of disasters to minimize the management and messaging required to just manage the team. The second is a forwarding logic required for supporting recipient hierarchy, in contrast to how ICN routers forward traffic based on the typical name hierarchy for content. For example, with the typical content hierarchy in (the COPSS enhancement to) NDN, when a recipient subscribes to a name (*e.g.*, `/sports`, indicating an interest in messages related to all sports topics) he would receive messages that are sent to the names under it in the hierarchy as well (*e.g.*, `/sports/football`, indicating the message is related to football under the sports category). Forwarding based on a longest prefix match enables this quite simply. However, with recipient hierarchies, a recipient subscribing to a name (*e.g.*, `/police/Aldgate`, for the policemen in Aldgate) should receive messages that are sent to the names above it in the hierarchy (*e.g.*, `/police`, for all policemen). This new forwarding logic should be added to the network to enable the efficient forwarding for disaster management.

Based on a careful study of different kinds of disasters (using sometimes limited publicly available information), we make the following specific contributions:

- We identify the requirements to an efficient communication platform in managing disasters;
- The design of a naming schema that can support both communication within the normal administrative hierarchy as well as supporting cross-jurisdiction/cross-functional communication in disaster situations;
- A plan-and-instantiate mechanism that allows the government to plan the “roles” (namespace structure) beforehand and dynamically instantiate the recovery plan for a disaster, thus minimizing the management and messaging required in the first stage of the disaster recovery;
- A new forwarding logic (recipient-hierarchy) to meet the need for communicating to dynamically formed groups sets of recipients for efficient multicast/anycast; and
- Qualitative and quantitative studies that show the benefit of CNS in terms of flexibility and efficiency in managing the disaster.

2. RELATED WORK

Here, we discuss the state of the art techniques that are, and can be used by emergency services.

2.1 Legacy and IP-based Emergency Services

Existing, legacy emergency service infrastructures that rely on circuit switched telephony (including mobile infrastructure such as MobileIP for voice calls) for emergency calls is not suited for data communication and cannot enable enhanced (interactive) services including video, written messages and contextual information. Recently, an attempt is being made to design and implement the *next generation emergency services* (NG112, NG9-1-1) [6–9] by adapting the IP infrastructure to meet the requirements of emergency services. A goal of this work is to allow citizens/authorities to contact emergency services with technologies they use to communicate every day. *E.g.*, work such as location-identification [10] deals with identifying location of a SIP caller, services such as LoST [8] map location to services based on service boundary [11] (*i.e.*, contact the closest or administratively correct police station). However, most of these work only focus on emergency calls from civilians (*e.g.*, [9]). Moreover, these design choices are affected by the limitations of the legacy as well as the IP infrastructure and cannot work out of the box in a fragmented scenario since they depend heavily on end-to-end communication.

2.2 Location Independent Architectures

ICN, as we stated before, shifts the focus from location dependent routing to forwarding messages based on the content identities. NDN [2, 12] is a popular variant of ICN. The current design of NDN adopts a URL-like scheme for content names, *e.g.*, this paper could have a name `/ICN16/CNS.pdf`. Content providers register the availability of content by its prefix. These prefixes are announced for global reachability in the Forwarding Information Base (FIB) of routers. Two kinds of packets are used: *Interest* and *Data*. An Interest is sent by a consumer to query for data. When forwarding an Interest, routers perform *longest-prefix matching* in FIB and find proper outgoing (inter)face towards the provider. Any data provider who receives the Interest can respond with a Data packet. Data packets follow the reverse path established by the Interest.

COPSS [5] extends NDN with push (or multicast) functionality. Instead of using ContentNames to identify con-

tent, COPSS uses hierarchical Content Descriptors (CDs) to describe content, *i.e.*, a content can have multiple CDs and a CD can identify multiple content items. CDs are also in the form of URLs, *e.g.*, the paper can have CDs */Networking/ICN*, */UniGöttingen/papers/CNS*, *etc.* A consumer interested in a CD can subscribe to the CD and will receive all the contents with the CD and its descendants. COPSS maintains a rendezvous point (RP) based subscription tree in a new data structure in the routers called Subscription Table (ST). As described in [13], it also allows the presence of multiple RPs to avoid traffic concentration.

The commonality across these solutions is that they propose the use of *name based forwarding* which avoids the early binding of forwarding messages to a specific location. CNS leverages the benefits provided by these solutions, while paying particular attention to the requirements of name-based communication in disaster situations to enhance their capability. Additionally, CNS can leverage LoST-like services to map roles to location dependent authorities.

2.3 ICN-DTN

Some early work has proposed the use of ICN in Delay Tolerate Networks (DTN)/fragmented environments. In [14], the authors present scenarios for using ICN in natural disasters. Work such as [15–18] deal with improving data delivery in DTN environments by leveraging benefits of ICN. [16] proposes a priority-based information dissemination/flooding mechanism in DTN, where the priority is based on the names; [17] proposes the use of ICN for vehicle to vehicle communication; [15, 18] propose to use ICN-based mules to spread information in DTN; [19] proposes an energy efficient message delivery mechanism that leverages collaborative communication in disasters.

CNS can make use of these solutions to perform data delivery in fragmented/DTN scenarios (disasters). Additionally, our work complements these efforts by providing a more comprehensive solution to handle the requirements of all kinds of disasters (adversarial, accidental and natural), both in fragmented and non-fragmented cases. Moreover, CNS details a naming mechanism that allows authorities to focus on the role instead of the individual that is performing that role at a particular point of time.

3. STUDY OF DISASTER SCENARIOS

We explore several example disaster situations to help us understand the requirements for the communication platform. According to the US National Protection Framework [20], disasters can be generally divided into 3 categories – adversarial, natural and accidental. We first look at adversarial disasters (with an example) which have the highest requirements on the communication platform and then extend our view to accidental and natural disasters to build a generic platform for all kinds of disasters.

3.1 Disaster Example – London Bombing 2005

At 8:50 am, July 7, 2005¹, a suicide bomb was detonated on the eastbound Circle Line train #204 traveling from Liverpool Street to Aldgate Station. Within 1 min, a second explosion took place on a Circle line train #206, going westbound from Edgware Road to Paddington. Approximately 2 min later, a third bomb was detonated on a southbound

Piccadilly Line train #311. At 9:47 am, a fourth bomb was detonated on the top deck of a #30 bus at Tavistock Square. The explosions resulted in 52 deaths, 700 people being physically injured and hundreds of people being directly affected.

The overall picture from 8:50 until about 9:15 was chaotic. Multiple, often conflicting, reports were being made, some to London Underground’s Network Control Centre, some to the emergency services, and some to the media. It was not clear what had happened, or indeed where. One major cause for the chaos was that the messages could not reach the right people (*e.g.*, first responders, resource managers).

Under the circumstances where the situation was not clear, first responders were unable to dispatch resources effectively. The first fire engines were dispatched to Praed Street instead of Edgware Road at 9:00 and it was not re-deployed until 9:37, nearly 40 minutes after the initial event. The survivors spoke repeatedly of the apparent lack of ambulances, equipment and supplies at the scenes, *even an hour or more after the explosions*, to the event review committee. Dispersal of patients to hospitals was also uneven because of a breakdown of communications within the Ambulance Service.

The lack of proper communication among different services hampered coordination. The London Emergency Services Procedure Manual clearly states that a “major incident” can be declared by any of the emergency services, the implication being that this will be done on behalf of all the services. However, different first responding services declared a major-incident separately, and at different time periods (*e.g.*, in Aldgate, the ambulance service realized it 19 minutes after the fire brigade did; in the case of King’s Cross, fire brigades were dispatched to the wrong sites and it is unclear when they realized that it was a site of a major incident). This late realization also affects the involvement, either directly or indirectly, of large numbers of people including other agencies such as the Local Authorities, National Health Service, Environment Agency, Military and Voluntary agencies. Certain hospitals in the vicinity were not aware of such an incident and did not participate in the rescue efforts for a long time.

Other departments/individuals can also play an important role in managing the disaster, which also needs proper communication with other disaster managers. *E.g.*, the London Underground Network Control Center put in an emergency services call to three sites (all correct places) at 8:59 am (only 9 min after the explosions). The records revealed that these calls did not result in the immediate dispatch of the emergency services to the scenes. For some reason, the message did not seem to get through to the right people. London Underground Emergency Response Unit – a small and little-known team which does not even have the right to use the blue light on the roads – played a crucial role in emergency response in the absence of the Fire Brigade at Russel Square. At Tavistock Square, there were no other ambulances at the scene at that time, but the bus was located outside the headquarters of the British Medical Association and doctors and other trained first-aiders came out of the building to care for the injured. If there was proper communication to the hospitals and nearby services, the victims could have been evacuated earlier.

All these issues demonstrate that the key to an effective response to a major or catastrophic incident is communication. This includes effective communication *within and between* the emergency, health, transport and *other services*,

¹This description of events is mainly based on [1].

Fig. 1: Content hierarchy vs. recipient hierarchy (hashed nodes: receivers, bold: name prefix in packet).

and also with the individuals caught up in the incident, and the public at large.

What happened in London on July 7, 2005 could happen in any country, city, at any time, as we have witnessed over the last decade. Anecdotal reports in newspapers indicate the outcomes in many of these cases were impacted by lack of timely and appropriate communication. Other than the adversarial disasters, governments also have to deal with disasters caused by accidents and nature, such as earthquakes, tsunamis, nuclear reactor incidents, *etc.* Although the cause and the scale of a disaster might differ, all major incidents can be expected to share some typical characteristics: 1) the involvement of numerous, different, agencies in the response, 2) the importance of effective communications within and between those agencies, and 3) the crucial importance of approaching each incident from the point of view of those directly caught up in it, either as members of public or as individuals involved in the response. The requirements on the communication platform in terms of dynamic group formation, and convenient role-based communication is similar in many of these cases.

3.2 Communication Platform Requirements

Based on studying several of these events, we arrive at the following common requirements for the communication in managing disasters:

- **Predefined roles & dynamically formed groups:**

Governments usually have plans to manage different disasters. However, these plans only consider the *roles* rather than the individuals/identities that “instantiate” these roles. This poses a challenge for current communication frameworks since they mainly focus on identities of individuals and reach them based on the communication devices they have. It is preferable to have a mapping from the role to the identity, with the mapping known to everyone authorized to manage and help the disaster. Therefore, the platform should be able to support communication based on such predefined roles and support a dynamic mapping from roles to identities without burdening individuals to manually maintain the mapping. Thus, when people are trying to provide information, they can reach the proper receivers more easily.

- **Efficient group communication:**

To deal with disasters at different scales, governments usually need to mobilize varying numbers of first responders. These responders should collectively be able to share information and receive commands. They might follow a separate control hierarchy, *e.g.*, a commander might want to send commands to the police team that is responding to the specific event, or even to all the first responders.

- **Content hierarchy vs. recipient hierarchy:**

The semantics for the recipient hierarchy, which is often used in the command chain, is different from what is used with content hierarchy in name oriented network architectures, such as NDN/COPSS. With content hierarchy, if a

consumer sends an Interest to `/Sports/Football` (the bold node in Fig. 1a), the interest will be sent to providers serving (by propagated FIB entries) `/Sports` and `/Sports/Football` (dotted nodes in Fig. 1a) according to the longest prefix match rule. Providers serving `/Sports/Football/Players` (or `Teams`) will not receive this request. Longest prefix match is sufficient to handle content hierarchy.

However, in a command chain (we refer to as recipient hierarchy), when a commander wants to send an interest to all policemen dealing with London Bombing (with name `/LondonBombing/Police`, the bold node in Fig. 1b), this message should reach policemen serving the FIB entries `/LondonBombing/Police` and its descendants like `Aldgate`, `Bishopsgate` (dotted nodes in Fig. 1b). This is not achievable using longest prefix match based forwarding. A similar challenge arises in pub/sub as well, when forwarding based on the subscription table entries. Therefore, we see a need to have a recipient hierarchy in the network in order to send packets to people serving names that are descendants of the prefix contained in the packet.

- **Priority-based communication:**

Prioritization is important in disaster communication. Extreme solutions like ACCess OverLoad Control (ACCOLC) do allow authorities to have reasonable communication, but at the cost of blocking all civilian traffic and at the risk of causing public panic. Therefore, when civilian communication is still desired in such situations, the system should not place a blanket block of all civilian communication (as is often the case with current telephony-based solutions). Instead, it should prioritize the communication among authorities, enabling an efficient command chain for managing the situation.

4. ARCHITECTURAL DESIGN

This section first provides an overview of CNS and then focusses on design details such as naming, use of templates and forwarding.

4.1 Architecture Overview

CNS holistically considers communication both for normal situations and for disaster scenarios, with a goal of using the same common infrastructure at all times. To find a suitable naming schema for CNS, we studied the use of flat names [3, 21], hierarchical names [5, 12] or even more complicated namespaces [22, 23]. We observe that organizational hierarchy is common, well understood and is often efficient in managing human interactions. The hierarchical structure is widely adopted in many situations, including managing first responder services, the military, *etc.* In our architecture, we try to represent what is already used in real world communications, so that the users do not need to change the organization or behavior they are already used to. We also want to have the network exploit the namespace to replicate and distribute the information efficiently to the group of recipients defined by the names. Therefore, CNS adopts hierarchical names for communication in disasters.

Fig. 2 shows an example namespace at a country level (UK). For communication among authorities under normal circumstances, CNS uses a naming structure to represent the administrative/organizational hierarchy (left side in the figure). For dynamically-created and possibly transient teams dealing with different incidents, a place holder `/UK/Incidents` is created (right side). Each incident (from a small incident

Fig. 2: Name Hierarchy in CNS.

such as a gas leak to a disaster such as a tsunami) will have a sub-namespace created under this place holder. Templates for different types of disasters can be planned beforehand and instantiated on seeing the disaster in order to minimize management and messaging overheads.

CNS is an *application-layer* design that helps the disaster managers to decide what names they need to use and how they are going to communicate when disaster strikes. It can run on any Information-Centric Network like [3, 12, 21–23] as long as there is a proper mapping from the hierarchical application-layer names to the identities used in the network. *E.g.*, each node in the hierarchy can have a GUID in MobilityFirst [3]. To send a message, the sender needs to carry all the related GUIDs. However, the GUIDs in MobilityFirst do not have relationships. To send a message to **FireBrigade** in the administrative hierarchy, the sender has to carry the GUIDs of all the descendants of **FireBrigade**, with associated traffic and computation overhead. CNS can also use assertion-based networks (*e.g.*, INS [22]) since it is easy to map the hierarchical structure to (XML-based) assertions. Nonetheless, each router in INS has to parse complicated XML queries before forwarding the information. While it is true that INS can provide similar functionality, solutions such as NDN provides sufficient functionality at a much lower cost. Therefore, we prefer to run CNS over NDN/COPSS like networks since they provide native support for hierarchies at a lower cost, so that CNS can exploit the network for efficient multicast/anycast, something that is critical for efficient disaster management.

4.2 Administrative Hierarchy

To provide convenient and efficient communication among people associated with various authorities as well as people outside these groups, CNS uses a name-based solution – a hierarchically structured name architecture to represent the organizational command chain. This is convenient as they only need to communicate with a “role” at the appropriate position rather than considering the individual who “instantiates” the role. As for mobility, since communication is based on names, individuals, including first responders will not have a new identity when they physically move from one network “location” to another.

Fig. 2 shows a possible namespace for communication for authorities in London, especially those responsible for safety, law and order, *etc.* All authority-related roles are under the prefix `.../Authorities`. We assume that there is central control (called the GOLD coordinating group in London) for all first responder services, but the division of responsi-

bility is different across departments. *I.e.*, the distribution of police stations is different from the distribution of ambulance pools, which is again different from the fire brigade stations. Therefore, the namespace assignment should follow the same organizational command chain as in the real world. The figure shows a possible division of responsibility for police services in London. Each sub-node can be further divided to correspond to the real-world command chain.

The communication model for normal circumstances is similar to what was proposed in [2] and [5]. The communicating parties either send unicast (similar to VoCCN [24]), multicast (similar to pub/sub in COPSS) or anycast (similar to query/response in NDN). Following Fig. 2, when a commander wants to send a message (multicast) to all firemen in London, he can simply send a message with name `.../London/FirstResponders/FireBrigade`. All the fire fighters listening (subscribed) to FireBrigade at WoodStreet, Bishopsgate, SnowHill, *etc.* would receive the message. With this functionality, each message will only be sent once and is replicated in the network. If a commander wants to talk to any firemen in WoodStreet, he can simply initiate VoCCN with callee=`.../FireBrigade/WoodStreet`. All the firemen listening on the channel are serving this FIB (or its prefix) and they will receive the Interest. Bi-directional communication can start after the basic handshake.

To deal with the duality between recipient hierarchies and content hierarchies, CNS modifies the forwarding engine to support the new semantics, and adds an extra bit in the packet header indicating which forwarding strategy routers should use (see §4.4). Namespace `112operators` is used to receive emergency calls from civilians. It can be further divided based on the requirement of the emergency call bureau. The network does not place any limit on civilian emergency calls, but in NDN, each data would have a signature from the data provider to enable identification of callers.

Namespace `.../London/Government` could be established for civilians to receive news from the government. The government can use this channel to broadcast alerts on disasters, report progress on rescue efforts, *etc.*

With name-based communication, CNS allows emergency management authorities to build up a hierarchy according to their real-world command chain. Officers in different departments can listen to (serve a prefix or subscribe) these channels to “instantiate” these roles. When communicating, authorities only need to communicate to the role rather than to the individual who is currently in that role. CNS can provide convenience, flexibility and efficiency for communication among authorities even in non-disaster situations.

Fig. 3: Template for bomb incident management.

4.3 Incident Response Hierarchy

In this sub-section, we will walk through how a government should prepare a disaster template beforehand, how an authority instantiates a new namespace in the hierarchy when it knows of a disaster, and how first responders listen to the new namespace based on their duty assignment and communicate with each other. Disaster relief of the London bombing will be used as an example.

4.3.1 Disaster Templates

While it is true that first responders can still use the channels established for normal operations, temporary command chains (different from the original organizational structure) may need to be set up based on the magnitude of the disaster. As a standard operation procedure (SOP), government emergency management agencies prepare different plans for different kinds of disasters [25]. These plans usually focus more on the roles/functions (*e.g.*, communication/power restoration, civilian relocation [26]) in the aftermath of a disaster rather than the exact scale or location of the response team, so that each plan can deal with a certain kind of disaster. The assignment of the responders is performed during the disaster based on the actual situation.

CNS allows emergency management agencies to prepare a namespace template similar to their existing disaster management plans. Fig. 3 shows an example template for bomb incident management based on the plan written by authorities in Marietta, Georgia, USA [27]. The root of the template represents the management center. It will be renamed with the event identity when the template is installed into the existing namespace. Under the root, there are usually the services that are involved in incident management, *e.g.*, Ambulance, Police, Fire, *etc.* The sub-namespace under the departments can be set up based on the responsibility of each department. *E.g.*, in [27], the responsibilities of the Marietta police department after “actual bomb or explosive device detonated” include: triage the people on site, evacuate civilians, search for explosives, *etc.* The planner can setup sub-namespaces like Evacuation, Triage, Explosive Search accordingly. These functions can be further divided and sub-namespaces created as needed (*e.g.*, **PerimeterEstablishment** and **CommandPost** in the figure).

The disaster plan can also include Hospital in the namespace, although hospitals might not have a node in the original administrative and organizational hierarchy. However, when there is a disaster, the government would need hospitals to stand by, wait for notifications and report their status. The planner for disaster management can therefore provide a namespace to the hospitals. The namespace can also be sub-divided based on the functions (*e.g.*, by distance or speciality) in case the disaster management officers want to send different messages to different kinds of hospitals.

Sometimes, disaster management officials may want to send messages only to the manager/commander of a certain team rather than the whole group (*e.g.*, a fire fighter

might need to report the situation to the on-site police supervisor). To deal with this requirement, the planner can set a Commander under each level of management (see the bold nodes in Fig. 3). The fire fighter can now send messages to namespace `.../Police/Commander` and only the police supervisor(s) who listen to that name will receive the message.

4.3.2 Dynamic Group Formation

Once there is a template, the primary disaster management commander can easily “instantiate” the template in the namespace when a disaster occurs. In Fig. 2, there is an “Incidents” sub-namespace. According to a disaster response hierarchy, such a namespace can also be placed at a departmental, state or national level to deal with disasters of different scales. What the commander needs to do is to provide a name for the incident, and then “copy” the whole tree of the template to the Incidents namespace. *E.g.*, **AldgateBombing20050707** can be seen as an example of instantiating the bomb incident management plan (Fig. 3). Note that instantiating a template does not change anything in the network. The routers do not store the namespace (no extra state). Nor will new routes have to be created until the first responders listen to the roles. However, the namespace exists in the application layer, and the first responders would receive information about the name(s) they should listen to and the name(s) they should communicate with. The form of communication could range from query/response to any-cast to multicast. This design helps to reduce the substantial amount of control messages exchanged immediately after instantiating a template. The network would build up the FIB gradually when first responders listen to the names (by propagating FIB/ST). Also, the roles that do not have people responding to it (it can happen in many cases since the template might consider a more complicated situation and involve more people) will not have extra FIB entries. Therefore, planners can feel free to instantiate a template, and/or design a more detailed and complex template without worrying about more state (control overhead) being used.

4.3.3 Role Instantiation & Authorization

The right to send/receive messages in a certain namespace has to be authorized, similar to the capabilities in the real world. Let us look at how an event may play out. When a disaster occurs, the administrator/commander contacts the departments in the administrative/organizational hierarchy. The departments would dispatch units to deal with the incident. Similarly, in CNS, the incident commander has a key to instantiate the incident namespace. He can use the key to certify the departments that will be involved in the disaster management. The keys can be provided to the departments, by just using the department namespace. The departments can use the key to further certify first responders. The keys to the first responders may be given following the administrative hierarchy.

On receiving the new key to the disaster namespace, first responders can serve prefixes or subscribe to CDs accordingly. The network will build appropriate dissemination paths based on the routing strategy. In Fig. 2, FireEngine2 (under Bishopsgate in the administrative hierarchy) is mobilized for AldgateBombing. The commander can even add a whole department into the disaster management namespace (*e.g.*, SnowHill and all its units are added to deal with Aldgate bombing).

The benefit of the new namespace is effective and convenient communication. *E.g.*, when an incident commander

wants to send instructions to the policemen who are establishing a perimeter, some of the policemen might come from WoodStreet while others might come from Bishopsgate. If the original administrative command chain were used, the commander would have to send the instructions twice, and the policemen from WoodStreet that are dealing with other duties will also receive them. This is an unnecessary burden and overhead on the commander, the network and the first responders. Using the disaster namespace, the commander only needs to send the instructions once and they will only be disseminated to the appropriate officers.

The requirements for disaster management may change according to the nature of the disaster and the judgement of the management officers. CNS also allows the officers to dynamically add sub namespaces as needed and new first responders can join the new namespace to participate in the disaster management. The procedure is similar to adding a new disaster namespace and we omit the details here.

4.4 Supporting Recipient Hierarchies

With CNS, we observe that communication along the chain of command for both the administrative and incident response is based on the recipient hierarchy, which is quite different from the content hierarchy.

Although it is possible to support recipient hierarchies without modifying the forwarding logic in an NDN/COPSS router, such a solution can result in inefficiency in the network. *E.g.*, consider the case where the police commander in Fig. 1b needs to get multicast calls/messages meant to reach *all* of the members dealing with London Bombing (group identified by the name /LB), or to the subset who are the members of the police department (identified by /LB/Police), but he does not need to get messages sent to a specific individual (*e.g.*, /LB/Police/Aldgate/PoliceA). To avoid that with a content hierarchy, he has to create and listen to a new name /LB/Police/Commander. This is equivalent to using specific multicast “channels” without taking advantage of the hierarchy. Then, upper-layer commanders have to send copies of each message to multiple, separate channels. This not only results in more traffic in the network, it is also undesirable in the real world since it places a considerable burden on the commanders to have to keep track and send to each of the names in the namespace.

Because of these concerns, we propose an additional forwarding logic for routers. A flag in the packet header can be used to indicate if the packet should be forwarded based on this recipient hierarchy or the usual content hierarchy. On receiving such packets, based on whether it is unicast/anycast (an Interest packet) or multicast (a Publication packet), the router would choose to look into the FIB or ST accordingly (following the COPSS design). The new logic requires the router to find a match in FIB/ST and forward packets to *any/all* entries under the name that is a match. *E.g.*, [28] provides a mechanism to perform efficient lookups in a data structure like the FIB and ST. On detecting a match, the router would only have to get the outgoing faces of any or all the descendants in the tree, and forward the packet accordingly. *E.g.*, on receiving an Interest with name .../FireBrigade and the special flag (an initial packet for a call to any responders in the fire brigade), the routers will forward it to a responder who serves .../FireBrigade/* (could be an engine group in SnowHill). On receiving a Publication with CD .../FireBrigade and the special flag

(a multicast message sent to all the officers in the fire brigade service), the routers will get all the outgoing faces of subscribers who subscribe to .../FireBrigade/*, and forward the packet accordingly. With this subtle change, we can support the new recipient hierarchy. The feasibility and efficiency of the modification is evaluated in §5.

4.5 Attribute-based Prioritization

During the course of a disaster event, citizens seek to communicate with one another to convey and enquire about their well-being and location. After the London bombing incidents reported by BBC, Vodafone experienced a 250% increase in the volume of calls and a doubling of the volume of text messages. Cable and Wireless handled 10 times the normal call volume of the Vodafone and O₂ networks. This sudden burst of traffic severely affected information exchange related to the incident, impacting: 1) communication between different services, which usually have to go through the civilian network; 2) communication between civilians and first responders for information updates; and 3) civilian calling for help, even among each other. Although the ACCess OverLoad Control (ACCOLC) feature adopted in telephony allows authorities with special devices to communicate on the civilian channel (by suppressing all other civilian traffic), using it would have caused even more severe issues: 1) authorities without special devices would not be able to communicate; and 2) public panic since civilians would lose communication at all. Therefore, a system based on logical prioritization (rather than blocking) is desirable, to allow the disaster-related communication to be guaranteed while the traffic among civilians is supported in a best-effort manner.

In NDN, a straightforward way is to use name-based prioritization. *E.g.*, in Fig. 2, we can create a rule to prioritize all the traffic with name .../Authorities and .../Incidents. However, this solution has a significant drawback – it complicates the namespace. Authorities might end up having a whole name hierarchy for prioritized traffic and the same (or very similar one) for non-prioritized traffic since for the same set of receivers (destinations), the priority can vary based on content. With multi-level prioritization, it would result in having a hierarchy for each priority level, with more states in the network in both FIB and ST. Even worse, first responders have to listen (subscribe/propagate FIB) to multiple names, which will place additional burden on both the network and the responders.

To decouple prioritization from the destination (receiver) of a message, we use an attribute field in the packet to indicate its priority level. Routers prioritize packets based on this attribute field. This attribute can be added/suggested by applications automatically. To prevent people from abusing the prioritized attributes, we can use a signature to validate if the sender can prioritize the message (and decide the priority level)². For communication among authorities, we can verify if the (key of the) sender is signed (directly or indirectly) by the key of authorities in the chain of trust. Of course, authorities can also send non-prioritized messages even with the same public key and the same destination. When a civilian wants to provide important information to authorities, he can also add the prioritized attribute. While it is true that the validation of keys and signatures would

²Signature can be added to both Data and Interest packets [29] for authentication.

Fig. 4: Evaluation result of CNS (w/ and w/o recipient hierarchy) in lab testbed.

Fig. 5: Messages transformation Haiti $\rightarrow$ San Francisco, on space and time.

cause overhead, we argue that it is an inevitable overhead either in the network layer or in the application layer. We believe that in-network validation is a more appropriate choice since it can prevent malicious content from even entering the network (compared to the forward and eventually-discard mechanism used by application-layer solutions). Efficiency-wise, we can make access points and gateways (or some network functions) perform the validation rather than having every router perform it, so that forwarding in the network does not suffer from the overhead.

The attribute field can also be used to alter the forwarding rules for the packets. *E.g.*, we can use an attribute to represent a civilian calling for help from nearby users. Routers can broadcast this message within a limited scope of a few hops, but with priority.

5. EVALUATION

With the help of a prototype of CNS deployed in our lab testbed and a synthetic data trace, we demonstrate the feasibility of implementing and deploying CNS and show the efficiency of the proposed recipient hierarchy. A real-world topology with a real-world trace is studied using our simulator (widely used in previous work [5, 13]) for a comparison between CNS and MobileIP [30], the current state of the art for communication based on location of individuals rather than dynamically created groups. We omitted some detailed data regarding the setting of parameters here (please see the extended material [31] for a detailed description).

5.1 Lab Testbed Evaluation

We first evaluate the feasibility and efficiency of our proposed recipient hierarchy (§ 4.4) in our testbed.

5.1.1 Data Set

Our lab testbed comprises 6 physical machines that are used as routers and are connected by links with a 100Mbps bandwidth and 10ms delay. We use a single server to emulate 63 users and these users can dynamically link (with 50Mbps bandwidth and 5ms delay) to any of the six routers, based on their movement pattern. The time interval between movements of a user is probabilistic, uniformly distributed between 2s to 120s. The users form a 3-level quad tree recipient hierarchy, and each node in the tree has 3 users.

The total simulation duration is 70 minutes including 4,390 reconnections in total. The total # of messages exchanged is 6,300 and the size of each message ranges from 1-99 packets (1,500 bytes per packet). Each message is assigned to a user (sender) and each sender can choose to multicast the message either to his own department (a node he is listening to) or to a subordinate department (a child node). Messages are sent based on a uniform distribution over a total period of 60 minutes (first message is sent at minute 5).

5.1.2 Evaluation Results

In Fig. 4, we compare the two variants of CNS (w/ and w/o recipient hierarchy) in terms of the total traffic, delay and # of packet lost, for varying # of messages. Message queueing at a router's busy outgoing face, can adversely impact delivery delay, especially at the RPs. Message loss is mainly caused by mobility. Nodes cannot receive messages till the network state (in FIB/ST) is properly established.

Our results show that CNS with recipient hierarchy has: 1) lower total traffic since only one copy of the message is sent from the sender while solution without recipient hierarchy has to send the same data to each leaf node in the hierarchy separately (see Fig. 4a); 2) lower delay since fewer packets go through the RP, thereby facing a smaller queuing delay (see Fig. 4b); and 3) lower packet loss rate since the solution is able to better aggregate the subscriptions (see Fig. 4c). Based on packets captured by Wireshark, we see that the computation overhead for forwarding each packet is <2%. The results show that the recipient hierarchy is a cost-effective enhancement to ICN for information dissemination. However, note that this result does not mean that content hierarchy is not efficient for the purpose it was designed. The main take-away is that with a proper namespace structure, and the new forwarding rule (recipient hierarchy), CNS is better suited for the application scenario of disaster management. Essentially, both approaches are needed for a complete communication framework.

5.2 Trace Driven Simulation

In order to perform a realistic evaluation of a disaster scenario, we needed a data set that consists of: 1) workload (messages sent and received), and 2) communication infrastructure and mobility pattern in the disaster. Due to the

Fig. 6: Simulation result of content hierarchy (Note the difference in scale for Load and Delay).

Fig. 7: Simulation result of recipient hierarchy (Note the difference in scale for Load and Delay).

lack of such a complete data set, we combined 2 real-world data sets to emulate a city-level disaster, if that were to happen, say *e.g.*, in San-Francisco. We first make use of the emergency messages that were sent in the aftermath of the Haiti earthquake [32,33] as the workload during a disaster. We then used a San-Francisco topology [34] with cab movement [35] to represent the communication environment first responders could face during a disaster. We realize that this is limiting, as mobility patterns are very likely to be different, but it does demonstrate the effectiveness of CNS.

The topology (see [31] for detail) consists of 232 routers distributed across 5 overlapping ISPs, and with each ISP having 5 non-overlapping domains. With MobileIP, we set up one home agent for each ISP at its root node. In the case of CNS, we have just one rendezvous point (RP) in the whole network. We used relatively low bandwidth links (100Mbps per inter-ISP link) to reflect the use primarily for emergency data. Router processing matches the link rate, while Home-Agent processing for redirecting packets is 2ms and location modification is 5ms. Although the data set is not large, our solution can scale according to other real world needs.

Haiti’s message data set consisted of only a subset of the actual messages sent, *i.e.* it had 3,131 messages of the total that were sent in the first month after the earthquake. Therefore, in order to scale it up, we compressed the time for sending the messages to 1 hour, to emulate the situation of a large number of messages being sent in the immediate aftermath of a disaster. The CDF of messages *vs.* time is shown in Fig. 5c. The CDF of messages used in our evaluation look the same (but with different time scale of 1 hour) and therefore omitted in the figure. The # of packets (1500 bytes each) per message were obtained by dividing the message size by 50. The messages in the data set have the latitude and longitude of their origin (see Fig. 5a, each dot represents a message and the darker color means many messages are sent at almost the same place). We mapped these messages into the San Francisco map by performing a linear transformation and scaling it (see Fig. 5b). Since the dataset is small, we can imagine this to be a small-scale disaster (*e.g.*, highway bomb). Each emergency message is sent by an end-host linked to a router nearest to the messages’ origin (*i.e.* based on the message’s latitude and longitude).

The movement of receivers is also based on the San Francisco cab data trace [35] on 2008/5/28 (Wednesday), with the long trace showing a similar pattern on a daily basis.

There were 494 cabs in action on this day. Since our message dataset is compacted to 1 hour, we divided the 1 day cab movement into 24 sub-traces to study the effects of an emergency (*e.g.*, a bomb detonation) at different time periods to correspond to different movement patterns and load.

5.2.1 Communication using Content Hierarchy

The Haiti message dataset is already separated into a hierarchy consisting of 8 major categories (*e.g.*, urgencies, urgency logistics, public health, *etc.*) with each category consisting of several sub-categories. The total # of categories is 36, and each message can belong to multiple categories. We configure each receiver to listen to 1 category, *i.e.* the cabs are seen as first responders moving around to help people and they could receive requests to answer emergency calls.

We compare CNS to MobileIP (Fig. 6) in terms of the # of message losses, aggregate network load and latency. Our results illustrate that CNS has a significantly lower loss rate (Fig. 6a), lower network load (Fig. 6b) and lower delay (Fig. 6c) as compared to MobileIP, regardless of which hour the disaster occurs. This is due to the fact that MobileIP relies on unicast, which results in more traffic, higher latency (path stretch), and congestion at the 5 home-agents. Moreover, in the case of MobileIP, the end-host cannot fetch the content from a neighbour who also received that data. Note that MobileIP consumes 6 times the network load compared to CNS and causes unacceptable delay (>200s).

On the other hand, CNS can aggregate subscriptions, lower the probability of message loss (when you are near an authority who also subscribed to the channel or a higher-level channel, a subscription can succeed within 1 hop) and ensure that there is less congestion on the RP even though there is only one RP as compared to MobileIP’s 5 home-agents.

5.2.2 Communication with Recipient Hierarchies

Finally, we use the Haiti message dataset as a basis to emulate the information exchange among authorities to study the benefits of the proposed recipient hierarchy as compared to using the content hierarchy approach. We build a recipient hierarchy containing 6 levels and assign the messages and cabs into the recipient hierarchy randomly (see [31] for detailed subscription relationship).

In Fig. 7, we compare CNS *vs.* the case without recipient hierarchy using the same metrics that we used to analyze content hierarchy. We observe that CNS using the recipient hierarchy has a slightly lower loss rate due to improved ag-

gregation. More importantly, CNS with recipient hierarchy outperforms the other variant in terms of aggregate network load (almost by half) and latency (by up to 80%) due to its improved design. While we used just a single RP to highlight the benefit of the proposed CNS with recipient hierarchy, solutions such as those proposed in [13] can be used to increase the number of RPs and load balance among them in order to handle higher load. When the traditional NDN-content query approach is used, we can achieve comparable performance as CNS using the recipient hierarchy when users (first responders) know exactly when and from whom a message is being sent so that they can send an interest using NDN. Otherwise, without the benefit of aggregation, such an approach only performs as well as unicast. These results confirm the lab testbed based results of §5.1.

6. CONCLUSION

We have proposed an approach using ICN to provide flexible and timely communication during and after a disaster. We identified the requirements for such an architecture by performing an extensive study of official reports and anecdotal reports in the aftermath of several actual disasters (including terrorist attacks). Our proposed architecture includes enhancements to the current ICN approaches for communication among authorities, especially for dynamically formed teams of first responders. A key contribution is the dynamic creation and evolution of incident related (sub) namespaces, recipient hierarchies, to represent the context and roles of first responders assigned to the disaster. A new enhanced forwarding strategy to support such recipient hierarchies is very useful to minimize the amount of message transmissions. With the help of a prototype and large scale trace-driven simulations, we highlight the quantitative benefits of CNS in terms of network load and latency as compared to an IP based solution.

We believe it is important to shift the focus on disaster communication from being an afterthought to being a first class citizen, exploiting emerging network architectures. Effective, convenient and timely communication could result in better outcomes, including fewer casualties.

7. ACKNOWLEDGEMENTS

This work is supported by the ICN2020 Project (Advancing ICN towards real-world deployment through research, innovative applications, and global scale experimentation), a research project supported jointly by the European Commission under its HORIZON 2020 (Grant Agreement No. 723014) and the National Institute of Information and Communications Technology (NICT) in Japan (Contract No. 184); and US NSF under Grant No. CNS-1455815. We thank our shepherd, Lixia Zhang, for her support and insightful feedback.

8. REFERENCES

- [1] L. Assembly *et al.*, *Report of the 7 July review committee*. Greater London Authority, 2006.
- [2] V. Jacobson *et al.*, “Networking Named Content,” in *CoNext*, 2009.
- [3] A. Venkataramani *et al.*, “MobilityFirst: A Mobility-Centric and Trustworthy Internet Architecture,” *SIGCOMM CCR*, pp. 74–80, 2014.
- [4] A. Anand *et al.*, “XIA: An Architecture for an Evolvable and Trustworthy Internet,” in *HotNets*, 2011.
- [5] J. Chen *et al.*, “COPSS: An Efficient Content Oriented Pub/Sub System,” in *ANCS*, 2011.
- [6] E. E. N. A. (EENA), “Ng112 project,” <http://www.eena.org/pages/ng-112>.
- [7] N. E. N. A. (NENA), “NG9-1-1 Project,” <http://www.nena.org/?NG911-Project>.
- [8] T. Hardie *et al.*, “LoST: A Location-to-Service Translation Protocol,” RFC 5222, Aug. 2008.
- [9] B. Rosen *et al.*, “Framework for Emergency Calling Using Internet Multimedia,” RFC 6443, Dec. 2011.
- [10] R. Barnes and M. Lepinski, “Using Imprecise Location for Emergency Context Resolution,” IETF Draft draft-ietf-ecrit-rough-loc-05.txt, Jul. 2012.
- [11] K. Wolf, “Location-to-Service Translation (LoST) Service List Boundary Extension,” RFC 6197, Apr. 2011.
- [12] L. Zhang *et al.*, “Named Data Networking (NDN) Project,” *Relatório Técnico NDN-0001, Xerox Palo Alto Research Center-PARC*, 2010.
- [13] J. Chen *et al.*, “G-COPSS: A Content Centric Communication Infrastructure for Gaming,” in *ICDCS*, 2012.
- [14] J. Seedorf *et al.*, “Using ICN in disaster scenarios,” IETF Draft draft-seedorf-icn-disaster-04, Oct. 2015.
- [15] E. Monticelli *et al.*, “An Information Centric Approach for Communications in Disaster Situations,” in *LANMAN*, 2014.
- [16] I. Psaras *et al.*, “Name-Based Replication Priorities in Disaster Cases,” in *NOM*, 2014.
- [17] L. Wang *et al.*, “Data Naming in Vehicle-to-Vehicle Communications,” in *NOMEN*, 2012.
- [18] A. Tagami *et al.*, “Name-based Push/Pull Message Dissemination for Disaster Message Board,” in *LANMAN*, 2016.
- [19] S. Kim *et al.*, “Power-Saving NDN-Based Message Delivery Based on Collaborative Communication in Disasters,” in *LANMAN*, 2015.
- [20] “National Protection Framework,” <http://www.fema.gov/media-library/assets/documents/97350>, 2014.
- [21] B. Ahlgren *et al.*, “Design Considerations for a Network of Information,” in *ReArch*, 2008.
- [22] W. Adjie-Winoto *et al.*, “The Design and Implementation of An Intentional Naming System,” *SIGOPS*, pp. 186–201, 1999.
- [23] W. Fenner *et al.*, “XTreeNet: Scalable Overlay Networks for XML Content Dissemination and Querying,” in *WCW*, 2005.
- [24] V. Jacobson *et al.*, “VoCCN: Voice-over Content-Centric Networks,” in *ReArch*, 2009.
- [25] U. of Florida, *The Disaster Handbook – National Edition*, ch. 3.7 The Role of Government in a Disaster.
- [26] F. E. M. A. (FEMA), “Continuity of Operations Planning Template for Federal Departments/Agencies,” Sep. 2013.
- [27] Marietta, Georgia, USA, “Marietta Police Department Bomb Incident Management Plan.”
- [28] Y. Wang *et al.*, “Wire Speed Name Lookup: A GPU-based Approach,” in *NSDI*, 2013.
- [29] J. Burke *et al.*, “Securing Instrumented Environments over Content-Centric Networking: The Case of Lighting Control and NDN,” in *NOM*, 2013.
- [30] C. E. Perkins, “Mobile IP,” *Communications Magazine*, vol. 35, no. 5, pp. 84–99, 1997.
- [31] J. Chen *et al.*, “CNS: Content-oriented Notification Service for Managing Disasters (Extended Material),” University of Göttingen, Germany. <https://www.net.informatik.uni-goettingen.de/publications/1947/PDF>, Technical Report IFI-TB-2015-04, Nov. 2015.
- [32] Wikipedia, “Timeline of relief efforts after the 2010 Haiti earthquake,” http://en.wikipedia.org/wiki/Timeline_of_relief_efforts_after_the_2010_Haiti_earthquake.
- [33] datahub, “Haiti Crisis Map,” <http://datahub.io/dataset/ushahidi/resource/81d058a8-173a-49d9-8ce9-4edf5e7cafc9>.
- [34] F. Zhang *et al.*, “EdgeBuffer: Caching and Prefetching Content at the Edge in the MobilityFirst Future Internet Architecture,” in *WoWMoM*, 2015.
- [35] M. Piorkowski *et al.*, “Dataset of Mobility Traces of Taxi Cabs in San Francisco, USA,” <http://crawdad.org/epfl/mobility/20090224/>, Feb. 2009.