Name-Based Replication Priorities in Disaster Cases

Ioannis Psaras*, Lorenzo Saino*, Mayutan Arumaithurai†, K.K. Ramakrishnan‡ and George Pavlou*
*Dept. of Electronic and Electrical Engineering
University College London, UK
Email: (i.psaras, l.saino, g.pavlou)@ucl.ac.uk
† University of Goettingen, Germany
Email: mayutan.arumaithurai@cs.uni-goettingen.de
‡ Rutgers University, U.S.A.
Email: kkramakrishnan@yahoo.com

***Abstract*—In the immediate aftermath of a natural disaster, network infrastructure is likely to have suffered severe damages that challenge normal communications. In addition to that, traffic substantially increases as a result of people attempting to get in touch with friends, relatives or the rescue teams. To address such requirements of a challenged network, we propose a communication framework based on messages that exploits *name-based replication* of content and enables ad-hoc communications with *spatial and temporal scoping and prioritisation* of named messages. We evaluate our design against less sophisticated replication strategies and show that important messages (*e.g.*, from first responders) get disseminated to more nodes than less important messages.**

I. Introduction

The most common effect of natural disasters on communication networks has been reported to be network fragmentation due to, for example, base-station failure, cut-off links, flooding or power outages [1]. Still, however, the communication network is the main system on which people rely in order to seek for help, distribute important information, manage rescue teams, or get in touch with friends and family. This inevitably increases the traffic demand on the mobile part of the network [2]. Recent reports from the Great East Japan Earthquake on the 11th of March 2011 have shown that people relied mainly on their mobile handhelds in order to communicate over the Internet with the outside world [3]. Traffic demand increased 9-fold, while at times, up to 64% of base-stations were out of order, resulting in roughly 90% of the calls being dropped [1], [2]. In this paper, we focus on the operation of mobile devices in case of disaster situations, that is, when networks get massively fragmented with parallel increase in traffic demand. We argue that in case of disasters *ad hoc* communication between mobile devices becomes a necessity, while efficient use of network and device resources is of crucial importance to foster communication between citizens and rescue teams.

We begin by identifying the special characteristics, as well as the requirements, of a disaster situation. We argue that in case of disasters information has to: *i) spread locally and for a limited amount of time*, and *ii) reach multiple destinations*. For example, information regarding trapped or wounded people has to reach any member that belongs to the rescue or first responders teams. No single person is responsible for rescuing people in need, since operations have to be managed by *several teams* in *specific areas* and in a *distributed manner*. Furthermore, important information from the first responders (*e.g.*, from fire brigade regarding shelter or food supplies) has to reach all interested parties within a specific area where help is (or will be) available.

Clearly, spreading information in disaster cases involves *space and time* factors that have to be taken into account as messages spread [4]. The standard mobile communications system does not support multi-user message dissemination, let alone the fact that much of its infrastructure might be out of action after a disaster (*e.g.*, base-station failure) and therefore, cannot be depended upon.

Infrastructureless communications have been studied heavily in the context of Delay-/Disruption-Tolerant Networking (DTN) [5], [6]. Several approaches from mules to robots, as well as specific message dissemination protocols [5] have been proposed to achieve message delivery in ad hoc networks. The vast majority of those solutions, however, target point-to-point message delivery, that is, delivery from a specific source to a specific destination. Furthermore, the operation of DTNs over the IP-based infrastructure of the Internet has led to much debate and finally design of a new architecture realised in the Bundle Protocol.

From the early days of its existence, the Information-Centric Networking (ICN) paradigm has made it clear that a *name-based communication model* can deal with infrastructureless networks more efficiently than the current, point-to-point, IP-based network [7]. More recent studies on the interface between ICN and DTN (*e.g.*, [8], [9]) have shown that name-based mobility is superior to IP-based mobile communications.

Based on the needs and requirements of disaster communications, in this paper, we propose NREP, a mobile *name-based replication* system [10], where message replication is limited by time and space, that is, within a certain geographic area and with specific life expectancy [4]. Last but not least, replication is optimised by prioritisation rules, *integrated within the information message's name* to favour spreading of the most important messages [6]. For example, we consider messages from first responders as more important than messages between friends. We focus on cases where the mobile network infrastructure is not available and therefore messages have to be *stored, carried and forwarded* by mobile devices. For the purpose of the present study we focus on information that has to be delivered to many recipients and ignore point-to-point applications, such as email.

Our findings show that prioritisation indeed achieves huge savings in terms of the dissemination of less important messages, in favour of dissemination of the most important ones, according to our prioritisation classes. By throttling the spread of less important messages, our proposed NREP scheme spreads important messages to nearly 100% of the population in all studied cases.

II. A NAME-BASED COMMUNICATION FRAMEWORK TO MEET THE CHALLENGES OF DISASTER SITUATIONS

A. Motivation

Current IP-based DTN networks focus on the destination of the content, *i.e.*, where the content should go. This reduces the flexibility of the network to make content-centric decisions such as if the data is worth transferring/replicating in the network. ICN, with its focus on content-centricity-based forwarding allows for nodes (*i.e.*, routers in case of the fixed Internet infrastructure) to make decisions based on the name of the content. Since nodes are aware of the content's or message's name, they can retrieve the data from their own cache or forward it towards one of the sources of the data. With a few exceptions ([11], [12], [8]), most of the on-going research [13], [14], [15] addresses non-DTN scenarios.

In this work, we attempt to leverage the benefits of ICN in the aftermaths of a disaster, where *ad hoc* DTN communication becomes essential in order to deal with fragmented networks and the increase in traffic demand. We argue for the need of a name-based forwarding/replication scheme, wherein intermediate nodes use a *Name* associated with each message to make decisions such as whether to replicate and if so, according to what priority, or otherwise, store(-and-carry) and for how long storage should be allocated. Moreover, we discuss the need to expose other parameters such as priority, time-to-live and geographical constraints *in the name* or as *attributes of the name*. This is done in order to help increase the efficiency of intermediate nodes to make decisions on storage and replication.

In the following we proceed to explain the design of our *name-based replication* framework, NREP.

B. Naming design

NREP borrows from ICN [7], [15] principles of using content names as the primary means for routing. However, unlike conventional ICN that is primarily designed to support *name-based routing* in an infrastructure-based environment, NREP is designed to operate in an infrastructureless environment and focuses on *name-based replication*, rather than routing. The design challenges of NREP are: *(i)* to identify what are the parameters that help differentiate between the various messages; *(ii)* to choose which of the parameters that influence message replication to include in the *name* and which to include as attributes; and *(iii)* to identify and understand the resulting trade-offs.

1) Parameters used for differentiation: Similar to the role played by a *name* in ICN, the *name* in NREP is also responsible for identifying the different types of content. Therefore, we recommend the use of a hierarchical namespace (as suggested in [7] and [15]) instead of a flat namespace, to allow nodes to filter content based on important parameters present in the longest prefix match. The hierarchical namespace should have a globally understood prioritisation value. For instance, the hierarchical name-prefix could look like: `Emergency/SOS` or `Emergency/Fire` where the former could be considered to have higher global priority than the latter.

Other parameters such as `user-defined-priority`, `space` [16] (the geographical reach within which the data is considered valid), `temporal-validity` [16] (lifetime of the content), `size` are also considered to be important input for filtering or prioritisation within each group identified by the name-prefix. For instance, specification that the content is only valid in `district/city/country` and/or is only valid till `temporal-validity` could be used to further prioritise among content that has the same name-prefix `Emergency/SOS`.

2) Where to place these parameters?: Though the namespace could be extended to have the parameters mentioned above, *e.g.*, `<user-defined-priority>/<temporal-validity>/<space>/<size>/../`, we believe that there is a need to keep these values as *attributes* instead of accommodating them in the *name*. These *attributes* could be set by the sender, the receiver or any authorised intermediate node to express additional information related to the content. Therefore, the *name* would only have prefixes that are globally known and based on which, nodes perform efficient longest prefix matches, whereas the *attributes* are those set by the sender/receiver and could be used for further filtering.

3) What are the trade-offs?: Advantages of such a distinction between placing some parameters as *attributes* instead of including them in the *name* are: *(i)* the namespace or prefix cannot be manipulated by individual senders/receivers; *(ii)* nodes with limited capacity can perform name-based filtering and forwarding whereas nodes with more resources (*e.g.*, base-stations or nodes provisioned and managed by first responders [17]) could in addition perform filtering by attributes; *(iii)* the senders/receivers could additionally assign their own desired priority for the content.

4) Advantage of the overall design choice: The combination of a globally understood *name*-based priority and the *attributes* such as `user-defined-priority`, `space`, `temporal-validity` helps optimise the decision making process for forwarding/replication of data based on the availability of buffer space, energy levels and duration of the interaction. This way, nodes have the opportunity to choose whether or not to send/receive content that satisfies a certain prefix. For instance, a node could decide to send/receive content whose name has a higher priority associated with it, and additionally if it contains specific attributes, such as `user-defined-priority/time-to-die/` where `user-defined-priority=high` and `time-to-die = "30th December 2013, 18:00 CET"`. If there is enough memory space for all `priority=high` data, then content selection could be performed with the following prefix: `user-defined-priority=high`.

Another advantage of the proposed scheme is that there is no need to exchange statistics/digests on every interaction between intermediate nodes that function as mules. Current DTN designs [5], [6] require that the mules look into the meta-data of the content, accumulate statistics and exchange this data when they come in contact with other mules. This is a cumbersome process, especially if the frequency of meeting other mules and the amount of data present in the network as a whole is high and dynamic. In contrast, according to NREP, we make a clear distinction between what is expressed in the *name* and what is expressed as *attributes*. In turn, mules with lower capability (either in terms of energy or memory space) can perform filtering and exchange data based only on the *name* associated with the message. For instance, two mules on meeting each other could initiate a transaction with a quick handshake until the receiving node's storage capacity is fully utilised or the energy restriction of the sender/receiver comes into effect. However, if one of these mules has higher capability, it can perform an additional filtering/sorting based on the *attributes*. For instance, one could start sending high-priority data, *i.e.*, data with prefix `priority=high`. Then, if space and/or energy permits [6], the nodes can start the transfer of data with prefix `priority=medium`.

C. Priorities and namespaces

As discussed earlier, the need for prioritisation in order to make efficient use of network resources and ensure that safety-critical messages get preferential access to network resources is of paramount importance in the aftermath of disasters. Safety-critical messages must be given higher priority over other low-priority traffic when they compete for the same network resources.

According to our initial design, the `name-prefix` is associated with a globally recognisable `priority` factor. For example, as shown in Table I, the NREP application is globally preset with the knowledge that the `SOS` name-prefix has higher priority than the `chat` name-prefix. Additionally, one could also envision an application where the client decides the priority of the message and assigns a priority value accordingly (*e.g.*, `user-defined-priority=High`). We could also envision a network where dedicated nodes look through the *attributes* and/or the content and set the `user-defined-priority` appropriately. The `temporal-validity` value can be represented as a time-to-die in absolute unix-time, *e.g.*, `1387414134` which implies that the content is valid till 2013-12-19T00:48:54. Similarly, the space value, *i.e.*, the area within which the data is valid can be represented by the following format `<type=circle;pos=x,y; radius=r>`, or `<type=rectangle, leftpos=x,y; height=h, breadth=b>`. Alternatively, the space value can be represented in the global map format, *e.g.*, `country/state/city/<postal-code>`. Below is an example list of priorities together with their characteristics in terms of space and time limitations.

1) High priority messages: Messages calling for help could use the name-prefix SOS (see Table I). Such messages have to spread quickly and should live long enough until help is received. In order to minimise misuse (selfish behaviour), messages sent with this name-prefix should be smaller in size and the time-to-live should not be very long. Otherwise, it will be difficult to stop the message from spreading even after help has been received. Moreover, a long expiry time could imply that too many people end up responding to it, thereby overutilising scarce resources that could be used somewhere else. If no response is received within the stipulated time, the client can increase the time-to-live and send the message again. A challenge associated with this name-prefix is to find a means to stop the dissemination, once a particular team has responded to it to avoid multiple teams responding to a single SOS call. To deal with this challenge, one could apply TTR-like techniques [18] between the members of the rescue teams, in order to better organise and manage operations.

Furthermore, messages from central state entities with instructions from first-responders (fire brigade, ambulance) to citizens need to spread to everyone and should not expire. Here, only the application residing on a limited number of authorised devices is allowed to send data with a suitable name-prefix such as `Government`, `Police` (see Table I). Messages notifying the arrival of rescue teams in an area at some fixed point in time to distribute first aid kits, water, food, etc. can be high priority too. People in the area should be informed and the message should be deleted after the rescue team has arrived.

2) Medium priority messages: Messages from individuals announcing the availability of food, water, etc. in a certain area should spread locally and be deleted after a period of time, as the resource will have been consumed. Similarly, messages on availability of shelter, electricity or communication capability available in an area should spread within that area, and need not expire since the shelter will be present for a long time. Such messages will have to be deleted only if conditions change, *e.g.*, shelter is full.

3) Low priority messages: Messages sent by individuals trying to get in touch with people in the area to get together and help each other are assigned lower priority. Such messages spread locally and normally can be deleted after delivery. These messages use the $Chat$ name-prefix as shown in Table I and therefore receive lower priority compared to more important messages.

D. Operation

Each message is labelled with at least three attributes:

- A priority level (explained in detail in Section II-C)

Name-prefix	Global-priority	Time-to-live	Space	Size	Sender Authorization	Recipient	Notes
SOS	High	Short	Closeby	Very-small	All	First responders	To use to ask for help
Government	High	Indefinite	All	Small	Officials	All	To inform all of food-shelter, danger
First-Responders	High	Indefinite	Depends	Small	First-Responders	All	To inform all of rescue-teams arrival
Warning	Medium	Indefinite	All	Very-small	All	First Responders	FR verify and publish to all
Police	High	Depends	Depends	small	Police	Police members	To chat among themselves
safe	Medium	Short	All	Small	All	Public/Family	To inform others that they are safe
chat	Low	Low	All	Small	All	Public	To chat among each other

TABLE I: Examples of name-prefix prioritisation and corresponding values for *temporal-validity, space and size*

- A spatial scope, *i.e.*, the geographical area outside which the message is no longer important.
- A temporal validity, *i.e.*, a timer at whose expiry the content of the message is no longer useful.

Each device stores received messages in its internal memory and keeps them as long as their timer expires and the device remains within the boundaries of their spatial scope.

Each time two or more devices are close to each other, they start exchanging the messages they are currently storing. Each device assigns a weight w to all the messages it holds and forwards them in decreasing order of w. This weight is calculated as a function of the distance from the origin of the message, the residual time validity and its priority.

$$w = \alpha f_d(d) + \beta f_t(t) + \gamma p \qquad (1)$$

where $f_d : \mathbb{R}^+ \rightarrow [0, 1]$ is a monotonically decreasing function of the distance from the origin of the message, $f_t : \mathbb{R}^+ \rightarrow [0, 1]$ is a monotonically decreasing function of the time elapsed since the message creation and $p \in [0, 1]$ is a value expressing the priority of the message and $\alpha, \beta, \gamma \in [0, 1]$ with $\alpha + \beta + \gamma = 1$.

A key challenge in priority-based replication is to decide whether to drop or assign a high-priority to a message that has already consumed a lot of resources and is therefore close to expiring or close to reaching the destination. This would be in contrast to a message that was just created and therefore has a high temporal validity and/or reach. Based on this decision f_d and f_t in Eq. 1 are either monotonically increasing or decreasing. Although we evaluate similar concepts in the next section, we leave a more elaborate investigation of this issue for future work.

Each mobile device may also decide whether or not to forward messages on the basis of its residual battery life. In fact, if battery life is scarce, a device may decide to only forward most important messages or no messages at all.

III. Performance Evaluation

A. Evaluation Setup

We evaluate the proposed framework in the ONE simulator [19]. We target scenarios where the memory capacity of mobile devices is limited in order for the name-based transfer prioritisation and replication to come into effect.

We loosely define six different example message classes, carry out extensive performance evaluations and present here two distinct scenarios. In our first scenario, we use baseline settings in order to highlight the importance of prioritisation (*i.e.*, time and space limits are the same for all nodes and all algorithms). In our second, more realistic disaster case, we apply different characteristics to different messages according to the message class they belong to. That is, for example, high-priority (HP) messages are set to higher TTLs, in order to inform as many users as possible, while low priority (LP) messages have shorter TTLs, as the information they carry will not be valid after long time periods (*e.g.*, chat messages).

We use two representative performance metrics, namely, the fraction of messages that keep spreading until their expiry and the average number of replications per message per class. According to the first metric, the longer a message stays (and spreads) in the network for, the higher the probability to inform more users. The second metric indirectly reflects the average number of nodes that receive each message per each of the priority classes. Again, the higher the number of replications, the more nodes informed by each message.

B. Scenario 1

We evaluate the baseline performance of different replication approaches, that is, without time and space limits. Instead, we vary the buffer size available to the mobile nodes, in order to see the effect of memory capacity. We present the fraction of messages that spread until their *"time of expiry"* for two (out of six) message classes.

We experiment in a 16 km^2 area where a total of 480 mobile nodes exchange messages. We use the default settings of the ONE simulator with both static and mobile nodes in the Helsinki city centre; the transmission range of nodes is 10m and the transmission speed is 250kbps.

The results are shown in Figs. 1a and 1b. Name-based replication (NREP) achieves considerably better performance for high-priority messages (Fig. 1a), while its performance drops for low-priority messages (Fig. 1b). Performance here is measured in terms of the fraction of messages that replicate until their expiry. As mentioned before, the longer a message stays in the network (*i.e.*, some node's memory), the higher the probability that it will inform more users. Instead, FIFO and RND appear to have similar performance in all cases. This is more clearly shown in Fig. 1e, where we present the number of replications per message per class for 5MB buffers. There, we see that inline with our design principles, NREP transmits more messages of higher priorities, while it leaves less space for messages of lower classes.

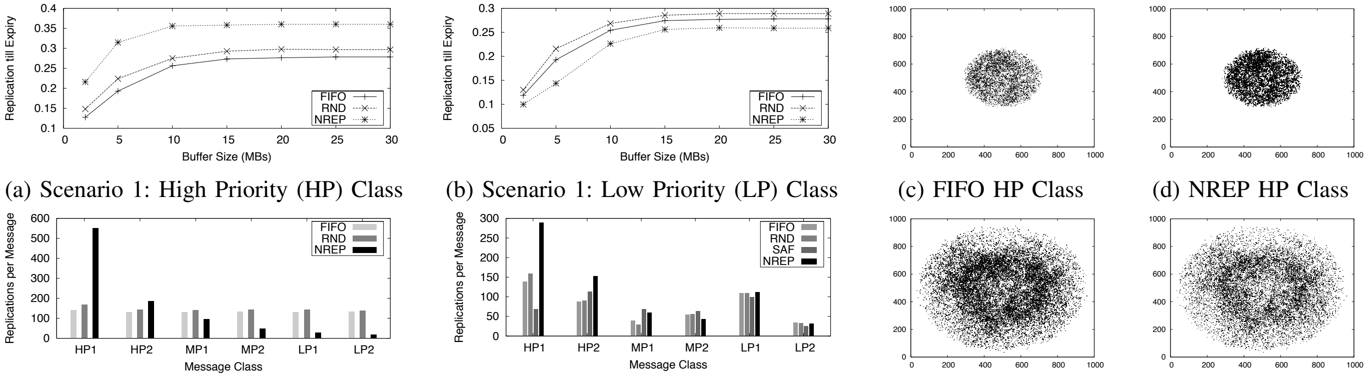

(a) Scenario 1: High Priority (HP) Class (b) Scenario 1: Low Priority (LP) Class (c) FIFO HP Class (d) NREP HP Class

(e) Scenario 1: Replications Per Message (f) Scenario 2: Replications Per Message (g) FIFO LP Class (h) NREP LP Class

Fig. 1: Results for Scenario 1 and 2. Figs. 1c, 1g, 1d, 1h capture the $1km^2$ simulation area.

It is worth noting that in Fig. 1e the replications for HP1 messages are higher than the number of participating nodes. This is due to the fact that some messages reach some nodes twice during their lifetime. Whether this is a desirable behaviour or not depends on the specific application in question, which we leave for future work.

C. Scenario 2

In our second scenario, we make use of two different Setups (Figs. 2a and 2b) to simulate a more realistic environment, where different message classes have different expiry times (TTL), generation intervals (GEN INT), generation co-ordinates (x, y) and replication areas (see Table II and Fig 2a). We assume 300 nodes that move according to the RandomWayPoint model within a $1km^2$ area. In Fig 2a, for instance, messages of classes HP1, LP1 and LP2 spread in the whole of the $1km^2$ area, while messages of classes HP2, MP1 and MP2 spread within a 300 meters radius (indicated by the larger grey circles in Fig. 2a). Furthermore, lower priority messages have lower TTL and are generated more frequently than higher priority messages (see Table II). In this scenario, we also evaluate the performance of the "Smaller Area First" (SAF) replication policy. According to this policy, messages with smaller space limits (*e.g.*, messages of classes HP2, MP1 and MP2 in Table II) are given replication priority over messages that spread to a larger area.

In Fig. 1f, we present the number of nodes each message has reached per message class. NREP reaches up to 288 nodes (out of 300 in total, ~95% of nodes) for HP1 - see Fig. 1f. NREP also outperforms the rest of the replication policies for HP2 messages, while it presents similar replication behaviour to the rest of the policies for the remaining message classes. Class-agnostic replication, as realised by FIFO and RND policies cannot provide preferential treatment to different messages, resulting therefore, in inferior performance. Note that the sum of messages per replication policy is not the same for all policies, as messages of different classes do not get generated with the same frequency and do not live for the same amount of time. That said, in terms of buffer occupancy over time, 100 messages of class x_1 might account for 70 messages of class x_2. NREP keeps more HP messages in the nodes' buffers and replicates more messages of these classes upon encounters with other nodes, than lower priority messages.

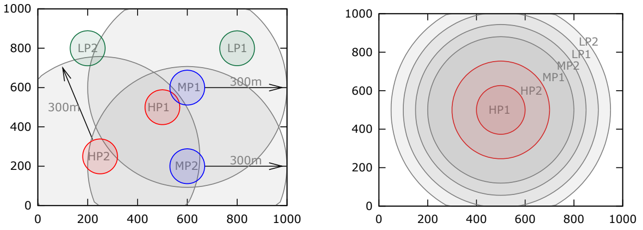

(a) Scenario 2 - Setup 1 (b) Scenario 2 - Setup 2

Fig. 2: Generation and Replication Areas of Message Classes. For Setup 1, small circles denote the areas where messages of corresponding classes are *generated*, whereas bigger circles (300m radius) denote areas where messages are *replicated*. Messages of classes HP1, LP1 and LP2 get replicated everywhere.

Finally, we see that SAF, as expected, does not perform well for messages that are set to replicate in large areas, *e.g.*, HP1 and LP1 and LP2. SAF is more efficient than the rest of the replication policies in case of small replication areas, *e.g.*, MP1, MP2. Further experimentation is needed to decide whether smaller or bigger areas should have higher priorities. As discussed in Section II-D this is a resource consumption issue, where one might argue that messages which have to travel further need to be prioritised over messages that have to cover smaller distances. In other words, prioritising smaller areas might starve messages that need to reach further out.

As a final evaluation step, we consider Setup 2 in Fig. 2b, where messages of all classes are generated in the middle of the area and also that the lower the priority of a message the further it spreads in the area (see replication circles in Fig. 2b). With this experiment we intend to show how messages get replicated in space, according to their priority. In Figs. 1c, 1d, 1g and 1h, we capture the x, y co-ordinates where replications take place within the $1km^2$ simulation area. We plot one high-priority (HP) and one low-priority (LP) class for FIFO and NREP, respectively. We see in Figs 1c and 1d that NREP HP messages are more densely replicated than FIFO ones. This results in less replications for LP messages of NREP compared to the denser replication of LP FIFO messages (see Figs 1g and 1h). This is inline to our design targets, according to which higher priority messages are favoured against lower priority ones.

Overall, we argue that smart replication priorities result in larger numbers of users becoming aware of important (*i.e.*, high priority) messages about the state of emergency. More sophisticated algorithms can surely be designed, but this study is a first step to unveil the potential of a *name-based replication* scheme with class prioritisation.

MSG Class	TTL	GEN INT	GEN x, y	REP Area (r)
$HP1$	115 mins	50 mins	500, 500	Everywhere
$HP2$	85 mins	50 mins	250, 250	300 meters
$MP1$	65 mins	35 mins	600, 600	300 meters
$MP2$	50 mins	35 mins	600, 200	300 meters
$LP1$	35 mins	15 mins	800, 800	Everywhere
$LP2$	15 mins	15 mins	200, 800	Everywhere

TABLE II: Message Class Values for Scenario 2

IV. Related Work

Disaster communications have recently concerned the research community which has worked towards improving the resiliency of communication networks during disasters. Studies in this area include the design of "Movable and Deployable Resource Units" [17], where the authors study how mobile communication units, *e.g.*, base stations can be deployed on the fly to replace damaged equipment. In [2], the authors design an information dissemination Social Networking Service (SNS), which is resilient to disasters and can improve performance in case of network disruption and fragmentation.

The authors in [18] design the *"Time To Return"* (TTR) routing protocol, which is used among the devices of first responders in order to carry information back to the headquarters as quickly as possible and without extensive replication, thus avoiding excessive resource and energy consumption. Although the works in [2] and [18] focus on issues related to our design space and target not only communication in case of disasters, but also energy efficient message dissemination, they do not cover applications where information has to be disseminated to multiple users. We argue that although such protocols and infrastructures are essential in case of disasters, a multi-recipient dissemination system is necessary in order to send and gather information from heavily affected areas.

In this respect, closer to our work is [10], where the authors build a Twitter application for the Android platform, which is also enhanced with a *"disaster mode"* operation. In disaster mode, the application is operating in an ad hoc manner and exchanges information with other nodes in the vicinity. The authors focus on the security considerations of such an infrastructure to enable authentication and prevent spam. It is important to stress that in case of a disaster and when infrastructure is not available in order to communicate with friends and family far from the disaster area, it is still important to make local communication available. People trapped in the same area will be much relieved to know that other people are near them and can help or provide resources and first aid. In our proposed name-based replication framework such communication is supported, while the framework is also able to distinguish between types of messages.

Last but not least, the concept of *Floating Content* has been recently proposed in [4] and [20] as a means for local message dissemination, in an infrastructureless manner to enable digital graffiti and social network applications between mobile devices without the need of a central server or Internet connectivity. The authors in [4] and [20] introduce the concept of message dissemination with time and space limitations, which is also central to our philosophy here.

The authors extend the concept of floating content in name-based communication environments in [20], but they do not consider prioritisation of message transfers. Therefore, when nodes meet they exchange messages of interest either in a FIFO or random manner, or based on the *anchor zone* where the message is being made available. As we have showed, this lack of message transfer prioritisation makes the work presented in [4] and [20] unsuitable for disaster situations, where the delivery delay might be of vital importance.

V. Summary and Conclusions

We have proposed a *name-based prioritisation and replication scheme* for messages in fragmented networks during disasters. Our scheme borrows ideas from the Floating Content [20] concept as well as from offline pub/sub systems that work in infrastructureless environments (*e.g.*, Twimight [10]), but enhances them in order to work in a name-based, Information-Centric environment, which provides benefits over IP-based, host-centric networks.

Our results show that indeed higher priority messages get disseminated to more nodes in the network, which might be of vital importance in case of disaster/emergency. Our proposed scheme does not take into account the energy of devices, which might be a scarce resource. Therefore, our immediate next step is to integrate energy considerations in the replication of messages and avoid transferring messages in excess when devices start running out of battery.

Acknowledgment

This research was funded by the joint EU FP7/NICT GreenICN Project, Contract No. 608518 and NICT No. 167.

References

[1] "Technical report on telecommunications and disaster mitigation," ITU-T FG-DR&NRR, Tech. Rep., 2013.

[2] T. Ogawara, Y. Kawahara, and T. Asami, "Information dissemination performance of a disaster-tolerant ndn-based distributed application in disrupted cellular networks," in *IEEE P2P 2013*, 2013, pp. 1–5.

[3] T. Sakaki, F. Toriumi, and Y. Matsuo, "Tweet trend analysis in an emergency situation," in *Proceedings of SWID '11*, pp. 3:1–3:8.

[4] J. Ott, E. Hyytia, P. Lassila, T. Vaegs, and J. Kangasharju, "Floating content: Information sharing in urban areas," in *IEEE PerCom*, 2011, pp. 136–146.

[5] M. Khabbaz, C. Assi, and W. Fawaz, "Disruption-tolerant networking: A comprehensive survey on recent developments and persisting challenges," 2012.

[6] I. Psaras, L. Wood, and R. Tafazolli, "Delay-/disruption-tolerant networking: State of the art and future challenges," in *Technical Report, University of Surrey, 2009*. [Online]. Available: http://www.ee.ucl.ac.uk/~uceeips/dtn-srv-ipsaras.pdf

[7] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *CoNEXT '09*.

[8] G. Tyson, N. Sastry, R. Cuevas, and et al, "A survey of mobility in information-centric networks," *Commun. ACM 2013*, pp. 90–98.

[9] X. Vasilakos, V. A. Siris, G. C. Polyzos, and M. Pomonis, "Proactive selective neighbor caching for enhancing mobility support in information-centric networks," in *ICN Workshop 2012*, pp. 61–66.

[10] T. Hossmann and et al, "Twitter in disaster mode: Security architecture," in *Proceedings of SWID '11*, pp. 7:1–7:8.

[11] L. Wang and et al, "Rapid traffic information dissemination using named data," in *Proceedings of NOM '12*, pp. 7–12.

[12] J. Wang, R. Wakikawa, and L. Zhang, "Dmnd: Collecting data from mobiles using named data," in *IEEE VNC 2010*, 2010, pp. 49–56.

[13] T. Koponen and et al, "A data-oriented (and beyond) network architecture," in *SIGCOMM '07*, pp. 181–192.

[14] L. Zhang and et al, "Named data networking (ndn) project," *Technical report NDN-0001, Xerox Palo Alto Research Center-PARC*, 2010.

[15] J. Chen, M. Arumaithurai, L. Jiao, X. Fu, and K. Ramakrishnan, "Copss: An efficient content oriented publish/subscribe system," in *ACM/IEEE ANCS 2011*, 2011, pp. 99–110.

[16] J. Kangasharju, J. Ott, and O. Karkulahti, "Floating content: Information availability in urban environments," in *IEEE PERCOM Workshops*, 2010, pp. 804–808.

[17] T. Sakano and et al, "Disaster-resilient networking: a new vision based on movable and deployable resource units," *Network, IEEE*, vol. 27, no. 4, pp. 40–46, 2013.

[18] A. Martín-Campillo, R. Martí, E. Yoneki, and J. Crowcroft, "Electronic triage tag and opportunistic networks in disasters," in *Proceedings of SWID '11*, pp. 6:1–6:10.

[19] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE Simulator for DTN Protocol Evaluation," in *SIMUTools '09*, 2009.

[20] J. Ott and J. Kangasharju, "Opportunistic content sharing applications," in *Proceedings of NOM'12*, pp. 19–24.