

Opportunistic networking overlays for ICT services in crisis management

Raffaele Bruno

IIT-CNR, Via G. Moruzzi, 1 – 56124 Pisa, I
raffaele.bruno@iit.cnr.it

Marco Conti

IIT-CNR, Via G. Moruzzi, 1 – 56124 Pisa, I
marco.conti@iit.cnr.it

Andrea Passarella

IIT-CNR, Via G. Moruzzi, 1 – 56124 Pisa, Italy
andrea.passarella@iit.cnr.it

ABSTRACT

ICT infrastructures are a critical asset in today's Information society. Legacy telecommunication systems easily collapse in the face of disruptions due to security incidents or natural disasters. Hence, there is an urgent demand for new architectures and technologies ensuring a more efficient and dependable support for various security missions, such as disaster relief initiatives, first responder operations, critical infrastructure protection, etc. In this paper we advocate the opportunistic networking paradigm to build a self-organizing overlay ICT infrastructure for supporting dependable crisis management services. Our opportunistic framework to “glues together” surviving parts of the pre-existing infrastructure with networks deployed on-demand and users devices, and supports dependable distribution of coherent, updated, and non-contradictory information distribution. Finally, to show the potential advantages of our solution, we present initial results on the performance of different types of opportunistic infrastructures, by particularly highlighting the gains of context-aware systems.

Keywords

First responders, opportunistic overlays, anycast communication services, group mobility.

1 INTRODUCTION

In today's modern society the creation, circulation and manipulation of information are activities that pervade many aspects of our cultural, economical and social life. Consequently, governments, economy and society in general, are becoming increasingly dependent on Information & Communication Technologies (ICT), which are the means of providing information. For these reasons, the communication infrastructures used to transport information are considered a critical asset of our society, such as the transportation and power supply infrastructures, and they should be protected and secured. In addition, the nature and extent of the threats jeopardizing our communications infrastructures are considerable higher today than in earlier times, because they are becoming “*unpredictable catastrophic events*” [1]. In these situations, the availability of a dependable ICT infrastructure is essential because most of the crisis management activities rely on the fast and dependable circulation of information between government entities, operators of critical infrastructures, and rescue teams (e.g., to organize rescue operations), as well as on the interaction of first responders with citizens and victims (e.g. to locate people, to distribute early warnings, etc.) [14]. However, the amount of information potentially generated in such situations is orders of magnitude higher than in normal operating conditions, and it should – very likely – be supported by a possibly severely damaged communication infrastructure. In addition, as demonstrated by recent natural disasters and security incidents (from 9/11 to Indian Ocean Tsunami) “*telecommunications was the greatest single area of concern*” [3] because current ICT infrastructures are not designed to withstand unplanned and unexpected disruptive events. All the reports [1,3,4,5] produced by either governmental or independent forums and committees to assess the causes of the communications breakdowns that have taken place in the aftermath of large-scale crises have pointed out that: i) rescue teams' private networks did not provide efficient support to the teams; ii) several teams were not able to communicate due to lack of interoperability between private networks; iii) public mobile 3G

This work was partially funded by the IST program of the European Commission under the HAGGLE (027918) FET-SAC project, and by the IIT-CNR Institute under the project “Networks for Critical Infrastructures”.

networks were severely overloaded; and iv) information delivered to first responders was incomplete, sometimes outdated and contradictory.

Our objective is to tackle the above challenges (discussed in more detail in Section 2) through an *opportunistic overlay self-organizing ICT infrastructure* that dynamically adapts to the network disruption level. As described in Section 3, opportunistic infrastructures consider disconnections, isolation of (set of) nodes and heterogeneous networks as a rule rather than an exception, and are thus designed to enable ICT services despite heterogeneity of underlying networking technologies, disruptions and disconnections of all kinds. When used in emergency scenarios, they provide several advantages with respect to conventional ICT infrastructure. For example, opportunistic solutions permit to exploit *all* available networking resources in a unique homogeneous framework, from surviving portions of standard ICT infrastructures (e.g., Internet, cellular networks) to on-demand networks deployed by emergency teams (e.g., mesh, vehicular, sensor networks), to individual users' mobile devices (PDAs, mobile phones, etc.). This allows for greater user reachability, relieves congestion on the most critical portions of the network. A more in-depth discussion of ICT services enabled by opportunistic overlay solutions is presented in Section 3.

Opportunistic solutions in general, and their application to emergency scenarios in particular, are very novel research areas, and lot of issues has still to be explored. In this paper (specifically, in Section 4) we present initial results showing that opportunistic infrastructures are a promising solution to provide ICT services during crisis management, even in *worst-case scenarios*, i.e., when conventional infrastructures are completely unavailable and only users' devices are used to transport messages.

2 CRISIS MANAGEMENT WITH DISRUPTED ICT INFRASTRUCTURE

In this section, we discuss the major technical challenges that nowadays hinder the deployment of a reliable, secure and efficient ICT infrastructure for crisis management and disaster relief, by focusing on recent emergency scenarios, such as Hurricane Katrina, 9/11, London bombings, etc.

First of all, we may observe that public safety agencies and law enforcement entities have long emphasized the need for dedicated wireless systems to efficiently support emergency response and public protection. However, a central lesson pointed out by recent crises is that private mobile radio systems maintained by public safety agencies, were outdated, incompatible [6,7] and difficult to rapidly deploy on the disaster area. Inefficiencies in the design or deployment of private networks are leading first responders and emergency managers to switch to public mobile networks. However, public mobile networks generally rely on dedicated infrastructures, adopt a centralized management of the communications resources, and exploit point-to-point links to interconnect the devices to other devices or control units. Thus, crucial system functionalities, as access control, connection establishment, support of mobility, etc. rely on the ICT infrastructure remaining almost intact. For these reasons, it is widely recognized that commercial systems are extremely vulnerable to disruptive events. The most recent security incidents and disasters have highlighted that even resource replication is not effective to ensure communications system resiliency because these backup solutions are often unable to handle the huge traffic volumes generated in the wake of a crisis situation.

There is now a general consensus that self-organizing architectures exploiting the ad hoc networking paradigm are the only available technological solution able to provide dependable communication services during crises [1,15] because they permit to quickly set-up autonomous "islands" of communication by creating self-organized peer-to-peer networks with mobile wireless devices. Figure 1 illustrates a typical crisis scenario affecting an urban environment where an incident has devastated the terrestrial infrastructure and various forms of self-organizing communication services have been deployed. For instance, rescue and public safety teams will bring in their own mobile devices able to communicate with each other without the need of any pre-deployed infrastructure [13]. Furthermore, it would be easy to install a mesh network in (part of) the security incident area to provide a wireless backbone for communications [11]. Ad hoc technologies [8,9] also enable to create self-organizing networks out of mobile devices (laptops, PDAs, smartphones, etc.) that people carry with them. The intrinsic re-configurability capabilities of ad hoc networks, and the use of multiple independent paths increases the availability and dependability of the wireless backbone through resilience to operational anomalies or security attacks. For instance, ad hoc networks may be able to connect to nodes in the surviving infrastructure and act as bridging components for isolated and disconnected fixed networks. However, the mobile ad hoc networking (MANET) concept cannot be seen as the "panacea" for the limits of legacy communications systems. In fact, MANET solutions assume the existence of continuous, rather stable end-to-end paths between communicating nodes, which is rarely true in disaster and emergency scenarios. Furthermore, because of the prohibitive environmental conditions or the large

scale of the security incident, it might be impractical to spread around a sufficient number of rescue vehicles so as to create well-connected mesh or vehicular networks. In this context it is more likely to envisage the case of clouds of connected handheld devices (e.g., palmtops carried by first responders) that will be just sporadically connected to each other, and, possibly, to the surviving part of the infrastructure. Thus, a novel approach should be devised to build a self-organizing network where different communications opportunities surviving or appearing in disaster areas (from surviving portions of traditional ICT infrastructures, to mesh, vehicular and sensor networks deployed on demand) may efficiently and transparently operate in combination, in order to guarantee continuity of emergency services, ensure critical data integrity and availability, and provide smooth degradation of communications services. In the following section we elaborate on how the *opportunistic networking approach*, a recent evolution of the traditional mobile ad hoc networking concept, can be exploited to develop a novel self-organizing ICT overlay infrastructure for supporting dependable communications services in emergency scenarios.



Figure 1: Required ICT infrastructure for crisis management after a disruptive event: opportunistic infrastructure complement surviving networks and enable communication services despite disruptions and disconnections.

3 OPPORTUNISTIC INFRASTRUCTURE FOR DEPENDABLE SERVICE PROVISIONING IN CRISES

The opportunistic (delay tolerant) networking paradigm is an evolution of the legacy MANET concept. Opportunistic network protocols enable end-to-end communication even when endpoints are *never* connected at the same time to the same network. More precisely, according to the *store-carry-and-forward* paradigm, intermediate devices store the messages when no forwarding opportunity towards the final destination(s) exists and exploit any future contact opportunity with other mobile devices to forward the messages [10,12].

Opportunistic networks are an outstanding opportunity to build more dependable ICT infrastructures in crisis-management scenarios. Specifically, our aim is to develop an *opportunistic overlay network*, which acts as a “glue” of all the heterogeneous communication resources that are available during crises and security incidents to the various actors involved in the emergency scenario. Such ICT infrastructure allows the combination of different technologies in a seamless and automated way, without requiring a full-functioning TCP/IP network as conventional P2P systems. Opportunistic overlays run on *any* communication resource available in the environment. In the most extreme case only wireless enabling technologies between couples of mobile devices are available. More in general, an opportunistic overlay provides ICT services also across portions of surviving infrastructure (TCP/IP fixed networks, cellular and satellite networks, etc.), mesh networks, vehicular networks, etc. These various network segments will be physically interconnected through special proxy nodes implementing gateway capabilities,

software compatibility layers, and multiple communications interfaces. In the simplest case, a gateway can simply be a multi-interface device that is able to communicate over more than one network segment with different link-level technologies. In most refined cases, the gateways can abstract the underlying networks, for the opportunistic overlay level, as a set of properties characterizing the behavior and capacity of the networks themselves. The main goal of the opportunistic infrastructure is *to exploit* the available technologies to provide communication services, given such gateways are available.

Also note that opportunistic overlays permit to cope with another important facet of heterogeneity, i.e., the heterogeneity of users' behaviors (e.g. in terms of movement patterns). Different behaviors may generate unstable topologies, creation and merging of partitions (because the network is also formed by devices carried by users). Opportunistic overlays are conceived to deal with these issues in the first place. We would like to stress the fact that different users' behavior is one of the most important facets of heterogeneity in crisis-management scenarios. While heterogeneity of technologies could be (at least partially) masked through common APIs and compatibility software layers, heterogeneity of users' behavior is something that does not depend on technology, and will be always present in such scenarios.

Opportunistic infrastructures of this kind are currently extensively analyzed in the framework of the European Haggie Project [18].

The advantages of this approach are manifold. First of all, current solutions require a unique stable infrastructure to allow communication between the different actors of an emergency scenario. Instead, an opportunistic overlay network provides a more efficient utilization of *all available* communication resources, enabling ICT services between all users of the crisis scenario, reducing congestion and increasing the number of people that have access to emergency services. The opportunistic ICT infrastructure we envisage is intrinsically dynamic and self-organizing, in the sense that it reconfigures based on the evolving network environment, and the heterogeneous users' behaviors (especially in terms of mobility patterns); and it transparently and organically grows as new devices, network segments or critical content become available (e.g., a new vehicular network set-up by rescue teams arriving on the disaster area). A further progress brought about by our solution is the increased dependability of the emergency services provided to first responders during security incidents, due to improved tolerance to frequent disconnections, network partitions and logical failures.

The communication services provided by opportunistic overlays could be exploited in a number of ways. Hereafter, we mention a few services for crisis-management scenarios that an opportunistic overlay enables.

Distributed information storage and retrieval. In a disrupted networking environment, end-to-end communications services should be complemented by data management techniques to ensure data availability and integrity. By leveraging an opportunistic overlay, algorithms could be designed to detect those nodes that have the highest importance for any given data type (e.g., nodes getting in touch very frequently with members of rescue teams will be important "repositories" for data of interest to the teams). As a special case, surviving elements of the infrastructure could be exploited as a sort of "proxy" that will either enable different clouds of users to get in touch with each other, or to temporarily store data until final destinations are able to retrieve them.

Decentralized service and resource discovery. An opportunistic overlay network will enable to design mechanisms to automatically and quickly discover, and adaptively control, the network resources and communication services that can be used to accomplish crisis response and management. For example, it could be possible to exploit structures induced by the users' behavior and movements, such as stable communication patterns or cliques between users (e.g., teams of first responders moving together in a coordinate fashion), to maximize efficiency and dependability of emergency services and crisis response.

Graceful performance adaptation. The opportunistic overlay network we envisage will be able to identify the level of degradation that impacted mission critical infrastructures. Then, it will be possible to autonomously activate the forms of communication services that are more appropriate to the operational conditions at the disaster area and the available communications resources. A further level of adaptivity will be also represented by the capability of the overlay network to adjust in a transparent way the performance level offered to emergency services. For example, full voice services may be downgraded to chat services with the degradation of the connectivity conditions.

3.1 Enabling communication services in opportunistic overlays

Opportunistic networking is still a relatively recent research area (the first papers on delay-tolerant networks have been published in 2003, e.g., [19]). Not surprisingly, one of the topics that are attracting researchers interest is

routing in opportunistic networks, i.e., how to efficiently enable end-to-end communication services between disconnected endpoints. Several routing schemes have been proposed for opportunistic networks (see [10,12] for detailed surveys on this area). It is possible to categorize them based on the amount of information they leverage to learn the features of the network they are operating in. In this paper we consider two opposite ends of the spectrum. On the one end, in pure dissemination schemes, nodes are oblivious to any available information. They just rely on aggressively spreading the messages in the network, seeking to reach the destination. On the opposite end of the spectrum, context-aware schemes leverage context information available in the network to selectively identify good next hops towards the destination. The most popular example of the former class is Epidemic forwarding (Epidemic for short) [20], that we also take as the reference point for this work. Epidemic adopts limited-scope, TTL-based flooding. When two nodes (say, A and B) get in touch, they exchange summary vectors that summarize the set of messages each one is carrying in its buffer. Then, node A (node B) receives from node B (node A) those messages that it is not carrying yet. For each received message the associated TTL counter is decreased. When the counter equals 0, the associated message can be only delivered directly to the destination. Note that nodes do not discard forwarded messages, and keep disseminating them upon encountering other nodes.

As representative of the context-aware schemes, we consider HiBOP, a solution fully specified in [21]. In HiBOP, the context is a collection of information that describes the environment in which the user moves, and the history of relationships among users. At each node, basic data used to build the context can be personal information about the user (e.g. name), about her residence (e.g. address), about her work (e.g. institution), etc. In HiBOP, nodes share their own data during contacts, and thus learn the context they are immersed in. Messages are forwarded following the store-carry-and-forward paradigm, through nodes that share more and more context data with the message destination. To estimate the match between an encountered node and the destination, HiBOP basically exploits i) context information about the users of the encountered node, and ii) context information seen by the encountered node in the recent past, on other nodes it met. Point ii) assumes that past encounters between users can be exploited to infer near future encounters. Even though over a rather short time frame, some degree of predictability can indeed be assumed also in emergency scenarios: the set of users (and, therefore, the context information) a user will meet in the near future is very likely to be correlated to the set of users met in the recent past. The main idea of HiBOP forwarding is thus looking for nodes that show increasing *match* with context information related to the destination. High match means high similarity between node's and destination's contexts and, therefore, high probability for those nodes to encounter each other.

Of course the advantage of context-aware protocols over dissemination-based protocols is a greater efficiency in terms of network resource usage. The drawback is the fact that, since context-aware protocols need to build correct statistics from context information, they require context information to circulate among nodes.

4 INITIAL EVALUATION OF OPPORTUNISTIC INFRASTRUCTURES

In this section we present selected initial results that highlight that opportunistic infrastructures in general, and context-aware opportunistic infrastructures in particular, are very promising solutions to enable ICT services in crisis-management scenarios. We firstly describe the simulation environment we use, and the types of services we consider in the evaluation.

4.1 Simulation environment and scenarios

In emergency scenarios, the high dynamism of users' behavior is one of the most critical challenges to design dependable communication services. Therefore, a correct characterization of the users' mobility patterns is fundamental to get realistic results. Generally, the various actors involved in a crisis situation do not move randomly, but the mobility patterns reflect the role and objective of each individual in the security incident area. More precisely, teams of operators generally work together on the same site and tend to move collectively and in a coordinated fashion. Similarly, in the aftermath of the disaster event, groups of people tend to stay together. We may also have single individuals that moves separately, for instance to link together different teams, or to provide support to groups of citizens. Finally, as a result of the evolution of the crisis environment, from time to time it may be necessary to suddenly relocate the rescue teams to deal with new critical situations. This corresponds to a collective movement of operators' teams that change their working site, and to a total reconfiguration of the disaster area.

To properly model the above describe mobility patterns we have adopted the Community-based Mobility Model, (CMM) recently proposed in [16], with the extensions to also model the relationship between users' movement and physical locations presented in [17]. According to this model, each node belongs to a community. Nodes that are in

the same community are called friends, while nodes in different communities are non-friends. Links between nodes represent relations among nodes, and links' weights represent the strength of the relations. Links towards non-friends result from relations across different communities. The model features two ways to represent these links. At the beginning of the simulation, for each node, each links towards a friend is rewired towards a non-friend with a probability equal to a model parameter called *rewiring* probability. Therefore, for each node, the sum of links' weights towards the group can be used to define a probability distribution. During the simulation nodes select the group towards which to move (with a uniformly distributed speed) according to this distribution: if s_{ij} is the sum of weights link between node i and group j , the probability of node i selecting group j for the next movement is given by $s_{ij} / \sum_j s_{ij}$. Alternatively, nodes in CMM can be instructed to always move towards the cell to which they are most attracted (i.e., the cell where they friends are). In this case, CMM also includes the notion of *travelers* that do not always move in the cell where they have more friends. From time to time, they move to the second most attractive cell (i.e., to the cell in which they have the second highest number of friends), and then get back to the most attractive cell afterwards. The effect of travelers and of the rewiring probability is exactly the same, i.e., representing relationships between different groups of users, thus accommodating for movements between different groups. Therefore, the two mechanisms are used interchangeably in our study. Note that these parameters also permit to factor in heterogeneity of users' behaviors, because they permit to describe users with different levels of interactions with other groups. Finally, in CMM, *periodic reconfiguration* occur, during which all groups change cell. During reconfigurations, collective movements of all nodes of any given group towards the target cell occur. Note that results provided in [16,17] show that CMM is able to reproduce the statistical features of real traces of humans' movements, collected within the Huggle project [18] (and publicly available on the CRAWDAD repository). Therefore, we use CMM in our evaluation as it has shown to be a valid and flexible tool to provide realistic synthetic movement traces.

In a crisis scenario, it is quite obvious that CMM communities represent rescue teams, firefighters' units, medical staff, groups of victims, etc., while a cell represents a working site, a medical camp, etc. Rewiring or travelers can be used to represent users that physically move between different groups to perform coordinated tasks requiring physical presence. Reconfigurations represent collective movements of squads or teams moving from one site to another in the crisis scene.

To evaluate the opportunistic ICT infrastructures we focus on ICT services and parameters that are particularly relevant for crisis-management scenarios. In a first set of experiments (Section 4.2) we consider a unicast messaging application and look at the performance of the opportunistic infrastructures under varying traffic loads. Scalability with traffic load is a key concern for crisis management scenarios, since the opportunistic overlay must be able to tolerate congestions and surges of network load without saturating the available communication resources¹. In a second set of experiments (Section 4.3) we evaluate the performance of the ICT infrastructures in anycast messaging applications, which are seen as one of the main communication paradigms in crisis scenarios. Specifically, we consider applications in which senders wish to communicate with *any* member of a target group, where groups are defined according to the CMM. Anycast services can be used, for example, by citizens seeking for rescue, to alert *any* available operator nearby. Or, they can be used by operators looking for *any* colleague with required special competences in a particular location of the disaster area. In Section 4.4, we consider the scenario of "closed" groups, i.e., groups in which *no* users are available to carry information from one group to another. This scenario represents, for example, groups of workers operating in different sites without interacting with each other, groups of injured people to be rescued from inaccessible places, etc. Clearly, understanding the performance of the opportunistic ICT infrastructures in such an extremely challenged scenario is a key target. Finally, in Section 4.5 we study the impact of different users' mobility patterns on unicast messaging applications. Specifically, we vary the reconfiguration parameter to study the effect on messaging applications of sudden encounter between nodes of different groups. To study a worst-case environment, in this case we do not include any links between nodes of different groups. Therefore, inter-group communication can occur only during reconfigurations, when nodes of different groups get in touch with each other because the trajectory of the collective movements of their groups intersect. Furthermore, we also vary the rewiring parameter, to understand how the performance of messaging applications depends on the level of inter-group movements. Note that this set of results provides indications on the opportunistic infrastructures

¹ This set of experiments provides initial results about the scalability properties of the opportunistic ICT infrastructure, although a comprehensive analysis of scalability features is not a target of this paper.

performance with respect to one of the most important aspects of heterogeneity, i.e., the heterogeneity of users' behaviors.

In the traffic-load and anycast experiments we consider 40 nodes divided into 8 groups randomly spread in a 5x5 grid (each cell being 250mx250m large). Each group is made up of 5 nodes. We consider one traveler per group, and a reconfiguration interval 9000s long. Nodes other than travelers have only social relations with their friends, i.e., inside their community. Thus, unless during reconfigurations, nodes move within their community, and only travelers are used to enable message exchanges between different communities. Upon each new movement, node's speed is sampled in the interval [2,9] m/s according to a uniform distribution. In our messaging application messages are supposed to carry a significant amount of information (e.g., a map sent by rescue teams with indication of evacuation paths, a short audio clip with information about people needing assistance in the crisis area, etc.). Therefore, messages' size is set to 50 KB, which is also consistent with typical message sizes considered, in general, in the opportunistic networking field² [22]. We randomly select 20 senders (uniformly among groups) at the beginning of each simulation run. Unless otherwise stated, the default interval between the generation of two consecutive messages at the same sender is modeled according to an exponential distribution, with average 300s. Message destination is a friend with 50% probability, and a non-friend with 50% probability. Among the friends and non-friends, the destination is chosen randomly according to a uniform distribution. In the anycast experiments the destination is always a non-friend. The group of non-friend destinations node is chosen randomly according to a uniform distribution. Finally, messages expire after (18000s), which is reasonable for delay-tolerant applications.

Even though our analysis is clearly not exhaustive, we consider this configuration as representative of crisis scenarios in which different teams have to intervene and communicate, without the availability of any surviving infrastructure. Therefore, this scenario also provides worst-case results for cases where some infrastructure is available after the crisis event.

The settings for the "closed"-group experiments, and the experiments with varying mobility patterns, have been chosen differently to better represent the addressed scenario, and will be described in Sections 4.4 and 4.5, respectively.

We define two sets of performance figures. The first one accounts for user QoS, in terms of average delay and message loss. In anycast experiments a message is lost if no member of the destination group receives it. In order to compare opportunistic solutions with different message loss rates also in terms of expected delay, a delay equal to the messages timeout value is considered for lost messages. The second set of performance figures measures the *resource utilization*, both in terms of average buffer occupation and bandwidth overhead. Buffer occupation and bandwidth overhead measure the congestion level of the system, and also help to explain the observed results in terms of delay and message loss. We measure buffer occupation as the average occupation of buffers during simulation runs. Bandwidth overhead is defined as the ratio between the total number of bytes exchanged over the network during a simulation run (also including traffic related to context management), and the total number of bytes generated by senders. Unless otherwise stated, hereafter we present confidence intervals (with 90% confidence levels) and average values of the performance figures, achieved by replicating 20 times, with independent seeds, simulation runs lasting for 90000s.

4.2 Scalability with the traffic load

In this set of experiments we vary the traffic load generated by the senders by considering an average inter-generation time between messages equal to 150, 300 and 600 seconds, respectively. Furthermore, we consider a limited maximum buffer size set to 50 messages. At light traffic load (average inter-generation time equal to 600s) the network is not congested, and the loss rate is very low (Figure 2). However, the more the load increases, the more the network becomes congested. Starting from medium loads (average inter-generation time equal to 300s) the loss rate of the two protocols tend to diverge, Epidemic experiencing a significant higher loss rate than HiBOP. This also results in the different trends of delay shown by Figure 3. At light loads Epidemic does not saturates network resources, and thus its flooding-like policy guarantees quicker message delivery with respect to HiBOP. This advantage becomes a drawback as the load increases, since Epidemic quickly saturates the resources (as shown by

² Note that, unlike in traditional IP networks, even such large messages are typically transferred in a unique *bundle* in opportunistic networks, as they should be self-contained in order to reduce the amount of round-trip interactions between communicating endpoints.

the high loss rates). At high loads, HiBOP is able to provide a *lower* expected delay thanks to its lower loss rate. The analysis of the resource consumption indices confirms this interpretation of the delay and message loss results. Specifically, Figure 4 shows that Epidemic quickly saturates the nodes' buffers, while with HiBOP buffers are seldom full (as the average buffer occupation is significantly lower than the buffer size). Finally, Figure 5 shows that the bandwidth overhead with HiBOP is always much lower than with Epidemic, thus confirming the higher efficiency of HiBOP in comparison with Epidemic (note that in Figure 5 the overhead decreases with the load, because a higher load implies a higher loss rate, and thus fewer messages surviving in the network with respect to the amount of messages generated). The authors of [21] have investigated the asymptotic limits in terms of resource usage, by considering unlimited buffers. They have shown that, on average, the consumption with HiBOP is one order of magnitude lower than the consumption with Epidemic.

This analysis indicates that a context-aware infrastructure is more efficient than a dissemination-based one in tolerating high traffic loads. Specifically, as the load increases, dissemination-based infrastructures quickly saturate the network resources. This leads to higher loss rate and expected delay with respect to context-aware infrastructures.

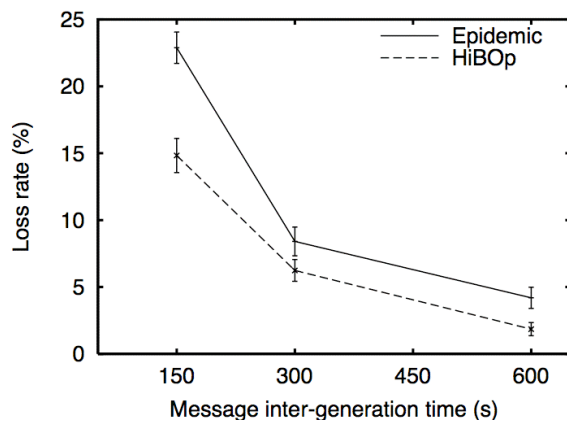


Figure 2. Message loss rate

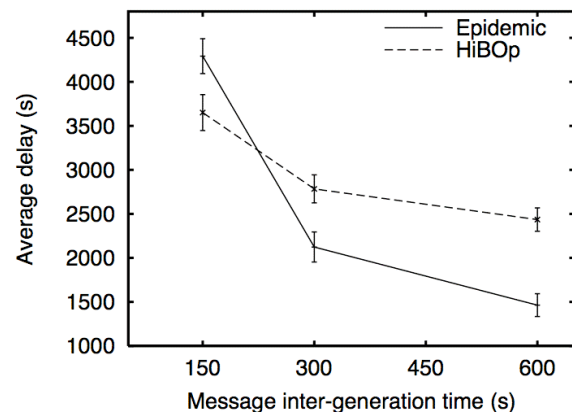


Figure 3. Average delay

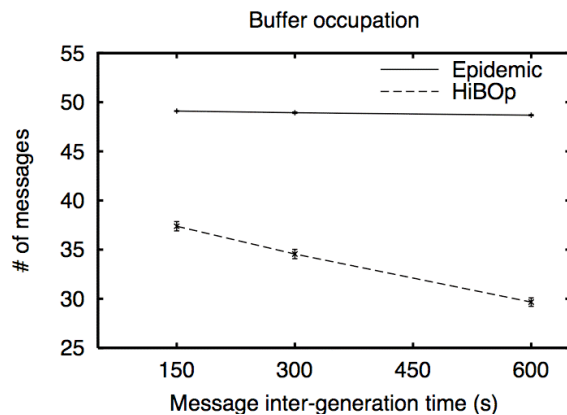


Figure 4. Buffer occupation

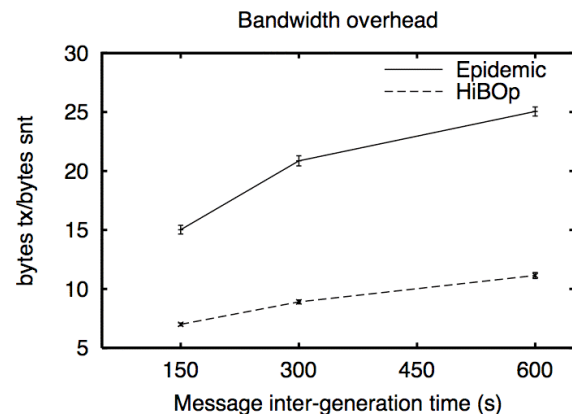


Figure 5. Bandwidth overhead

4.3 Anycast experiments

We show in this section results related to anycast applications, derived by varying the maximum buffer size at nodes. Specifically we consider unlimited buffers, and buffer size equal to 20 messages. Considering unlimited buffers allows us to show the asymptotic behavior of the two infrastructures.

Results are qualitatively similar to the ones presented in Section 4.2. In terms of message delay and loss rate, HiBOP becomes more efficient than Epidemic as the resources becomes more and more limited. Specifically, it achieves a

significantly lower message loss rate, and is also more efficient in terms of average delay (see Table 1 and Table 2). As expected, in terms of resource consumption Epidemic quickly saturates the buffers while HiBOP does not (see Figure 6). In the limit case when no buffer limitations are considered, the HiBOP buffer occupation is again one order of magnitude lower (note the logarithmic scale on in Figure 6). Finally, the bandwidth overhead is much lower with HiBOP than with Epidemic in either case (see Figure 7).

These results show that a context-aware infrastructure, such as HiBOP, is able to support also anycast (and, thus, group communication) services. Again, with respect to dissemination-based solutions, a context-aware infrastructure provides a much more efficient solution, which generates less resource congestion, and under realistic resource limitations, achieves lower loss rates and delay.

	Epidemic	HiBOP
Buff = 20	46.34±3.21	26.88±3.76
Buff = inf	0.08±0.16	4.00±0.94

Table 1. Message loss (anycast)

	Epidemic	HiBOP
Buff = 20	4.82±0.34	3.96±0.37
Buff = inf	1.36±0.12	2.85±0.16

Table 2. Average delay (anycast, 10³s)




Figure 6. Buffer occupation




Figure 7. Bandwidth overhead

4.4 Communication between nodes in isolated groups

In this set of experiments we consider “closed” groups without any traveler relaying information between groups. In this scenario, communication between groups occurs only when members of different groups happen to come close enough for their mobile devices to be inside each other’s radio range. It is important to understand how the frequency of such contacts impact on the performance of the opportunistic ICT infrastructures.

To model this scenario, we consider a 3x3 grid with 9 groups of 5 nodes each. Just one node, located in the upper left cell sends messages, destined to a node in the lower right cell. We configure the mobility model so that nodes do not move outside their groups. Thus, the only way a message can reach its final destination is through edge contacts with nodes of different groups. We simulate different contact frequencies between members of different groups by varying the devices’ transmission range: The higher the range, the higher the contact frequency. We use three values for the transmission range, i.e. 62.5m, 125m and 250m, representative of very low, medium and high frequency. We only show the asymptotic results with unlimited buffers, which provide a worst-case scenario for the context-aware infrastructure.

The bottomline of the results is that context-aware infrastructures are not suitable for networks in which context information cannot circulate due to too sporadic contacts between different groups. This is because, since the sender and the destination are in different groups, context information of the destination has to circulate towards the source's groups for HiBOP to be efficient. At very small transmission ranges (62.5m) HiBOP is not able to deliver acceptable QoS. HiBOP needs a minimum contact frequency between different groups to spread context information around. Indeed, in the 125m case HiBOP restores acceptable QoS at least in terms of loss rate, and is fully effective in the 250m case.

Also in this case Epidemic and HiBOP behave differently with respect to the bandwidth overhead (Figure 8). Epidemic overhead steadily increases with the contact frequency, as more opportunities to flood the network become available. At 62.5m HiBOP overhead is low because it seldom forwards any message. As context data is not circulating, all nodes in the sender's group almost equally suitable to carry the messages closer to the destination. At high transmission range the context data is circulating effectively, and therefore good paths can be identified soon. In intermediate cases there is a transitional regime in which HiBOP becomes effective in terms of delay and message loss at the cost of a high overhead. Note that Epidemic is not able to exploit rich connectivity scenarios (transmission range equal to 250m) without flooding the network.

These results suggest that a hybrid scheme can be the correct approach for networks with varying levels of context-information spread. When groups are very isolated, dissemination-based schemes seem the only way to enable ICT services between groups. As soon as context information spreads (a bit more) in the network, context-based routing becomes a preferable solution. An interesting follow-up of this work is how to exploit context information to distinguish these different scenarios and customize the operations of the ICT infrastructure accordingly.

	Range (m)	Epidemic	HiBOP
Loss rate (%)	62.5	0	65.79±9.29
	125	0	0
	250	0	0
Delay (s)	62.5	531.79±19.14	15579±734.45
	125	103.00±2.59	568.08±157.71
	250	23.35±0.52	1.51±0.64

Table 3. User perceived QoS

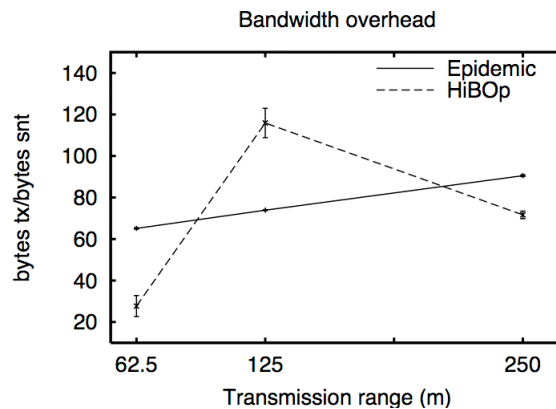


Figure 8. Bandwidth overhead

4.5 Sensitivity to varying users mobility patterns

In this section we analyze the impact of varying users' movement patterns on the communication performance for unicast messaging applications. We firstly focus on varying the reconfiguration parameter (i.e., the frequency of groups collective movements), then on the rewiring parameter (i.e., the probability of inter-group movements). To have a reasonable number of inter-group links (and, thus, to improve the statistical confidence of our results), we slightly modified the simulation setup, by considering three larger groups, each made up of 10 nodes. As in the case

of the anycast experiments, we consider here unlimited buffers, to characterize the asymptotic behavior of the infrastructures.

	Reconfiguration period (s)	Epidemic	HiBOP
Loss rate (%)	2250	0	0
	9000	5.52±1.46	8.16±1.68
	36000	24.12±1.31	25.64±1.30
Delay (s)	2250	907.10±67.08	1202.52±91.09
	9000	3204.58±278.70	3651.68±295.05
	36000	5445.11±161.53	5615.43±225.93

Table 4. User perceived QoS (reconfiguration)

It is worth recalling that, when varying the reconfiguration period, we set the rewiring probability to 0. Thus, except for reconfigurations, nodes do not have chances to meet. The reconfiguration period varies between 2250s, 9000s, and 36000s.

Table 4 shows the QoS performance as a function of the reconfiguration period. As expected, both packet loss and delay increase with this parameter, because messages addressed outside the group of the sender are forced to wait for a reconfiguration. Note that, even though HiBOP provides higher loss rate and delay, the difference with Epidemic is quite thin. These results clearly show that HiBOP is able to identify very good paths even during sporadic, sudden contacts during reconfigurations among nodes belonging to different groups. Again, the good performance in terms of user QoS shown by HiBOP comes along with a drastic reduction in resource usage. Figure 9 shows the average buffer occupation. HiBOP is much less greedy in spreading messages, and therefore the buffer occupation is drastically reduced. Finally, Figure 10 shows the bandwidth overhead of the two infrastructures. It allows us to highlight a main difference between HiBOP and Epidemic, related to how they react to movement patterns. Reducing the reconfiguration interval (from 36000s down to 2250s) means increasing the forwarding opportunities, because nodes get in touch with more peers more frequently. Epidemic does not use these additional "connectivity resources" wisely, as it is based on flooding. Therefore, the bandwidth overhead greatly increases. In the case of HiBOP, as nodes mix more and more (reconfiguration intervals equal to 9000s and 2250s), more overhead is generated, because more contacts become available, which may possibly lead to paths towards the destination. However, the rate of increase of the HiBOP's overhead is significantly lower than the one of Epidemic, thus showing a much more judicious use of the available network resources. These results confirm that, also in this case, the HiBOP infrastructure prevents network congestion by sparingly using the available network resources (without affecting the performance in terms of delay and loss rate).

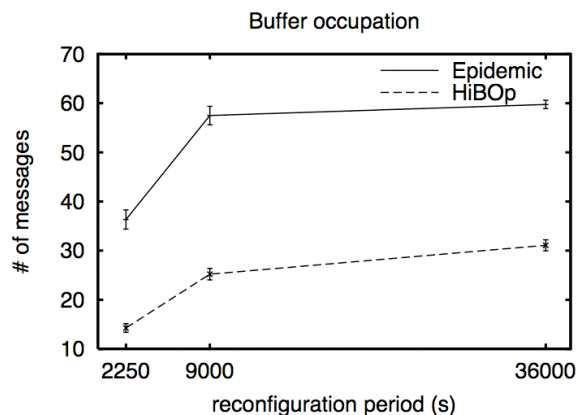


Figure 9. Buffer occupation

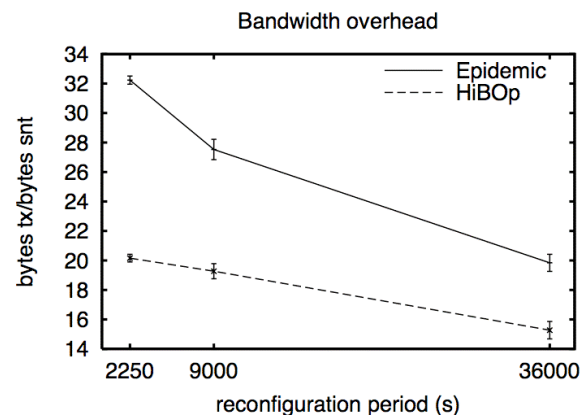


Figure 10. Bandwidth overhead

The last set of results we present show the sensitiveness of the compared infrastructures to varying level of mixing between groups, represented by varying the rewiring parameter. We consider three values for the rewiring probability, representing low, medium, and high connectivity between groups. As far as the QoS performance

figures (Table 5), the loss rate is negligible for both infrastructures, (so we do not show it), while – as expected – the average delay decreases as more users move between different groups (i.e., for increasing values of the rewiring parameter). The HiBOP's performance is still not far from the bound represented by Epidemic. In terms of resource consumption, again HiBOP shows to be much more effective than the Epidemic solution (see Figure 11 and Figure 12). When the rewiring probability increases, more nodes move between different groups, and thus more paths become available connecting nodes belonging to different groups. With both infrastructures, this results in lower buffer occupation, because messages arrive at the destination more quickly, and thus messages occupy buffer resources for lesser time. In terms of bandwidth overhead, it is interesting to note the *opposite trends* of the two infrastructures. At higher mixing rates, Epidemic overuses the additional resources that become available thus resulting in higher bandwidth overhead. On the contrary HiBOP leverages users' mixing (and the resulting spread of context information) to identify good paths more and more accurately. Thus, it needs fewer transmissions to carry the messages to the destination, thus resulting in *lower* bandwidth overhead.

	Rewiring probability	Epidemic	HiBOP
Delay (s)	0.03	130.28±20.59	170.86±25.86
	0.1	83.20±8.57	129.42±12.51
	0.5	73.69±7.16	104.91±8.87

Table 5. User perceived QoS (rewiring)

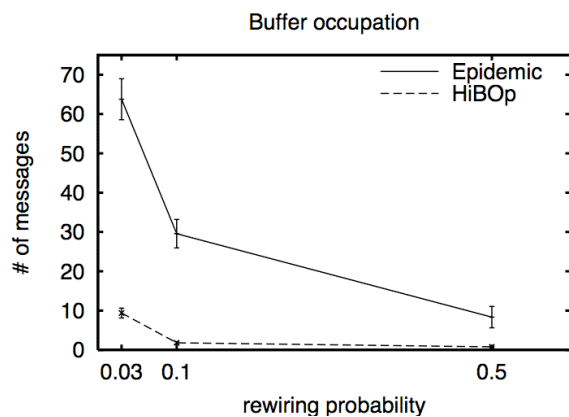


Figure 11. Buffer occupation




Figure 12. Bandwidth overhead

5 CONCLUSIONS AND FUTURE DIRECTIONS

In this paper we have proposed the opportunistic networking paradigm to build a dependable, dynamic and self-organizing overlay ICT infrastructure for crisis management. Unlike current state-of-the-art solutions, opportunistic overlays permit to exploit any network resource available in the crisis site, encompassing survived trunks of pre-existing infrastructures, operator networks, open emergency networks deployed on-site (e.g., mesh, vehicular networks), and even single mobile devices of people involved in the crisis scenario. We have presented an initial set of results showing that opportunistic overlays (and context-aware systems in particular) are a promising approach for crisis-management ICT infrastructures. Unless in particularly adverse conditions, context-based overlays actually provide an effective congestion control mechanism, and, with respect to dissemination-based overlays, provide acceptable QoS while greatly reducing resource congestion. Therefore, they provide a dependable ICT infrastructure that does not saturate network resources even under high traffic loads.

Several aspects of opportunistic networking for emergency scenarios are still to be explored. Among them, we believe that data management services are one of the most compelling and interesting research issues. Also, the characteristics of opportunistic overlays in terms of, e.g., scalability with respect to very large number of nodes should be extensively investigated. Another interesting research direction is the design of hybrid systems that automatically select the best type of opportunistic communication service (e.g., between dissemination-based and context-aware), based on the evolving configuration of the emergency network.

REFERENCES

1. ESARB (2006) Meeting the Challenge: the European Security Research Agenda - A report from the European Security Research Advisory Board.
2. European Commission (2006) FP7 Cooperation Work Programme - Theme 10: Security, *Call 1*.
3. London Regional Resilience Forum (2006) Looking Back, Moving Forward – The Multi-Agency Debrief.
4. Reid, J. and Jowell, T. (2006) Addressing Lessons from the Emergency Response to the 7 July 2005 London Bombings.
5. US Homeland Security (2006) Hurricane Katrina: A Nation Still Unprepared.
6. US National Task Force on Interoperability (2005) Why Can't We Talk?.
7. Hatfield, D. and Weiser, P. (2005) Towards A Next Generation Strategy – Learning from Katrina and Taking Advantage of New Technologies.
8. Conti, M. and Giordano, S. (2007) Multihop Ad Hoc Networking: the Theory, *IEEE Communications Magazine*, 45, 4, 78-86.
9. Conti, M. and Giordano, S. (2007) Multihop Ad Hoc Networking: the Reality, *IEEE Communications Magazine*, 45, 4, 88-95.
10. Pelusi, L. and Passarella, A. and Conti, M. (2006) Opportunistic Networking: Data Forwarding in Disconnected Mobile Ad Hoc Networks, *IEEE Communications Magazine*, 44, 11.
11. Bruno, R and Conti, M. and Gregori, E. (2005) Mesh Networks: Commodity Multi-hop Ad Hoc Networks, *IEEE Communications Magazine*, 43, 3, 123-131.
12. Z. Zhang, "Routing in Intermittently Connected Mobile Ad Hoc Networks and Delay Tolerant Networks: Overview and Challenges", *IEEE Communications Surveys*, vol.8, no.1, First Quarter 2006..
13. Rocchetti, M. and Gerla, M. and Palazzi, C.E. and Ferretti, S. and Pau, G. (2007) First Responders' Crystal Ball: How to Scry the Emergency from a Remote Vehicle, *Proceedings of the IEEE IPCCC 2007*, New Orleans, LA, USA.
14. Manoj, B.S. and Hubenko-Baker, A. (2007) Communication challenges in emergency response, *Communications of the ACM*, 50, 3, 51-53.
15. Dilmaghani, R.B. and Rao, R.R. (2007) Future Wireless Communication Infrastructure with Application to Emergency Scenarios, *Proceedings of the IEEE WoWMoM 2007*, Helsinki, Finland.
16. M. Musolesi, C. Mascolo, "Designing Mobility Models based on Social Network Theory", *ACM MC2R*, July 2007.
17. C. Boldrini, M. Conti, A. Passarella, "Users Mobility Models for Opportunistic Networks: the Role of Physical Locations", Proc. of IEEE WRECOM, Rome, Italy, July 2007.
18. "Haggle: A Novel Communication Paradigm for Autonomic Opportunistic Communication", European Commission FET-SAC IST 027918 Project (2006-2009), <http://www.haggleproject.org>
19. K. Fall, "A delay-tolerant network architecture for challenged internets", Proc. of *ACM SIGCOMM*, 2003.
20. A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks", Tech. Rep. CS-2000-06, CS Dept., Duke University, 2000.
21. C. Boldrini, M. Conti, I. Iacopini, and A. Passarella, "HiBOP: a History Based Routing Protocol for Opportunistic Networks", *Proc. of the IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM 2007)*, Helsinki, Finland, June 18-21, 2007.
22. J. Ott, "Application protocol design considerations for a mobile internet". In *Proceedings of First ACM/IEEE international Workshop on Mobility in the Evolving Internet Architecture* (San Francisco, California, December 01 - 01, 2006). *MobiArch '06*. ACM, New York, NY, 75-80