

---

## **Autonomic behaviour of opportunistic network routing**

---

Chiara Boldrini, Marco Conti  
and Andrea Passarella\*

CNR-IIT,  
Via G. Moruzzi,  
56124 Pisa, Italy  
E-mail: c.boldrini@iit.cnr.it  
E-mail: m.conti@iit.cnr.it  
E-mail: a.passarella@iit.cnr.it  
\*Corresponding author

**Abstract:** In opportunistic networks, end-to-end communication among users does not require a continuous end-to-end path between source and destination. Network protocols are designed to be extremely resilient to events such as long partitions, node disconnections, etc. which are very features of this type of self-organising *ad hoc* networks. This is achieved by temporarily storing messages at intermediate nodes, waiting for future opportunities to forward them towards the destination. Clearly, designing routing and forwarding schemes is one of the main challenges in this environment. In this article, we provide a survey of the main approaches to routing in purely infrastructure-less opportunistic networks, by classifying protocols based on the amount of context information they exploit. Then, we provide an extensive quantitative comparison between representatives of protocols that do not use any context information, and protocols that manage and exploit a rich set of context information. Mainly, we focus on the suitability of protocols to adapt to the dynamically changing network features, as resulting from the user movement patterns that are driven by their social behaviour. Our results show that context-aware routing is extremely adaptive to dynamic networking scenarios, and, with respect to protocols that do not use any context information, is able to provide similar performance in terms of delay and loss rate, by using just a small fraction of the network resources.

**Keywords:** autonomic communication; context-aware routing; delay-tolerant networks; opportunistic networks; routing in mobile networks; self-organising networks.

**Reference** to this paper should be made as follows: Boldrini, C., Conti, M. and Passarella, A. (xxxx) 'Autonomic behaviour of opportunistic network routing', *Int. J. Autonomous and Adaptive Communications Systems*, Vol. x, No. x, pp.xx-xx.

**Biographical notes:** Chiara Boldrini is a PhD student at the Institute for Informatics and Telematics of the National Research Council (CNR), Italy. She holds an MSc (2006) degree in Computer Engineering from the University of Pisa. She is working on socially-aware protocols and models for routing, data dissemination and mobility in opportunistic networks.

Marco Conti is a Research Director at IIT, an institute of the Italian National Research Council (CNR). He co-authored the book 'Metropolitan Area Networks' (1997) and is a co-editor of the books 'Mobile Ad Hoc Networking' (2004) and 'Mobile Ad Hoc Networks: From Theory to Reality' (2007).

He has published in journals and conference proceedings more than 200 research papers related to design, modelling, and performance evaluation of computer-network architectures and protocols. He served as general or programme chair for many conferences including: *IEEE MASS* 2007, *IEEE WoWMoM* 2005 and 2006, *IEEE PerCom* 2006, *ACM MobiHoc* 2006, *IFIP-TC6 Networking* 2002. He is an Associate Editor-in-Chief of *Pervasive and Mobile Computing Journal*.

Andrea Passarella is with IIT-CNR, Italy. Previously, he was a Researcher at the Computer Laboratory, Cambridge, UK. He holds a PhD in Computer Engineering (Pisa University, Italy). He is working on opportunistic self-organising networks, with special emphasis on p2p and routing protocols, mobility models and energy management. He serves in the Programme Committees, among others, of *IEEE WoWMoM*, *PerCom* and *MASS*. He was the TPC Co-Chair of *ACM/SIGMOBILE MobiOpp* 2007, Vice-Chair for *ACM REALMAN* (2005/2006) and *IEEE MDC* (2006). He is an Associate Technical Editor for *IEEE Communications Magazine*, and is in the Editorial Board of the Inderscience *IJAACS Journal*.

---

## 1 Introduction

The opportunistic networking idea stems from the critical review of the research field on Mobile *Ad hoc* Networks (MANET). After more than 10 years of research in the MANET field, this promising technology still has not massively entered the mass market. One of the main reasons of this is nowadays seen in the lack of a *practical* approach to the design of infrastructure-less multi-hop *ad hoc* networks (Conti and Giordano, 2007a,b). One of the main approaches of conventional MANET research is to design protocols that mask the features of mobile networks via the routing (and transport) layer, so as to expose to higher layers an Internet-like network abstraction. Wireless networks' peculiarities, such as mobility of users, disconnection of nodes, network partitions, links' instability, are seen – as in the legacy internet – as exceptions. This often results in the design of MANET network stacks that are significantly complex and unstable (Borgia et al., 2005).

Opportunistic networks (Pelusi, Passarella and Conti, 2006b) also aim at building networks out of mobile devices carried by people, possibly without relying on any pre-existing infrastructure. However, opportunistic networks look at mobility, disconnections, partitions, etc. as *features* of the networks rather than exceptions. Actually, mobility is exploited as a way to bridge disconnected 'clouds' of nodes and enable communication, rather than a drawback to be dealt with. More specifically, in opportunistic networking no assumption is made on the existence of a complete path between two nodes wishing to communicate. Source and destination nodes might never be connected to the same network, at the same time. Nevertheless, opportunistic networking techniques allow such nodes to exchange messages. By exploiting the *store-carry-and-forward* paradigm (Fall, 2003) intermediate nodes (between source and destination) store messages when no forwarding opportunity towards the final destination exists, and exploit any future contact opportunity with other mobile devices to bring the messages closer and closer to the destination. This approach to build self-organising infrastructure-less wireless networks turns out to be much more practical than the conventional MANET paradigm. Indeed, despite the fact that opportunistic network research is still in its early stages, the

opportunistic networking concept is nowadays exploited in a number of concrete applications (in Section 2, we provide a brief overview of them).

Until now, most of the research efforts in the opportunistic networking area focus on routing and forwarding issues<sup>1</sup>, due to the inherent complexity of the problem (Pelusi, Passarella and Conti, 2006b; Zhang, 2006). Therefore, after describing in more details the main concepts of opportunistic networks and their practical use cases in Section 2, in Section 3, we provide a brief survey of the main routing approaches available in the literature. Specifically, we categorise protocols based on the amount of *context information* they exploit, by identifying three main classes, i.e. context-oblivious, partially context-aware and fully context-aware protocols. The main idea behind using context information is to enable routing protocols to learn the network state, autonomously adapt to its dynamic evolution and thus optimise their operations. In the final part of the article (Section 4), we provide performance results to evaluate the suitability of this idea in real routing protocols. To replicate realistically the users' behaviour, we consider a mobility model (Home-cell Community-based Mobility Model; HCMM) that has shown to realistically reproduce real human movement patterns as driven by users' social relationships and social behaviour. We exploit the model's parameters to study how different routing approaches react to various levels of dynamism and users' sociability. We compare the performance of Epidemic Routing and History-Based Opportunistic Routing protocol (HiBOp), which are representatives of the opposite ends of the spectrum of possible approaches, i.e. context-oblivious and fully context-aware protocols, respectively (Section 4).

By analysing their sensitiveness with respect to a number of parameters, we show that context-aware schemes are able to provide similar levels of Quality of Service (QoS); in terms of message delay and loss rate), by spending a *small fraction* of the resources spent by context-oblivious protocols. Even more interestingly, we find that context-aware systems are much more suitable to autonomously learn the features of the network they are operating in, and the behaviour of users as determined by their social relationships. We show that, unlike context-oblivious systems, context-aware protocols are able to correctly adapt their operations accordingly. This results in a much more judicious use of the available resources, also when the network scenario abruptly changes. Finally, we draw conclusions and identify research directions in Section 5.

A preliminary comparison of Epidemic Routing and HiBOp has been presented in Boldrini, Conti and Passarella (2007b). In this article, we provide a detailed framework of the research efforts on routing in the opportunistic research area, and also provide additional performance results to better compare the different routing schemes.

## 2 The opportunistic networking concept and its applications

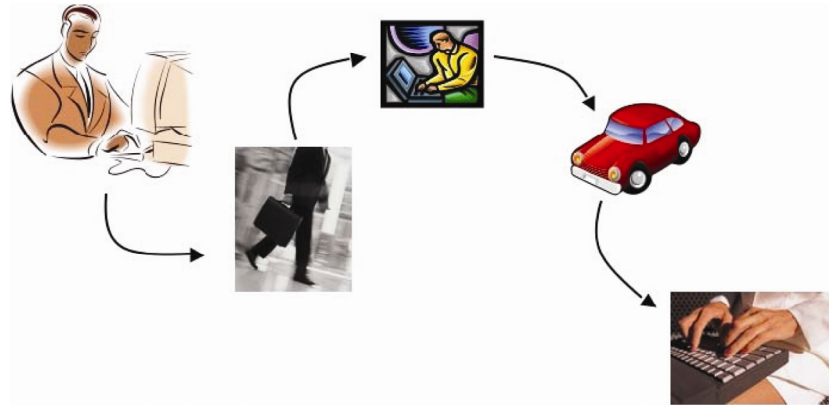
Opportunistic networks share several concepts with Delay Tolerant Networks (DTNs). The DTN architectures defined by the DTN IRTF Research Group (<http://www.dtnrg.org/docs/specs>) focus on a scenario in which independent internets, each characterised by internal Internet-like connectivity, are interconnected through a DTN overlay. In order to achieve end-to-end connectivity, the DTN *overlay* exploits occasional communication opportunities among the internets, which might either be

scheduled over time (e.g. due to the activation of a satellite link), or completely random. In general, in conventional DTNs the points of possible disconnections are known.

Opportunistic networks can be seen as a generalisation of DTNs. Specifically, in opportunistic networks no a-priori knowledge is assumed about the possible points of disconnections, nor the existence of separate Internet-like sub-networks is assumed. Opportunistic networks are formed by individual nodes, that are possibly disconnected for long time intervals, and that opportunistically exploit any contact with other nodes to forward messages. The routing approach between conventional DTNs and opportunistic networks is therefore quite different. Since in DTNs, the points of disconnections (and, sometime, the duration of disconnections) are known, routing can be performed along the same lines used for conventional Internet protocols, by simply considering the duration of the disconnections as an additional cost of the links (Jain, Fall and Patra, 2004). Since, opportunistic networks do not assume the same knowledge about the network evolution, routes are computed dynamically while the messages are being forwarded towards the destination. Each intermediate node evaluates the suitability of encountered nodes to be a good next hop towards the destination.

For example, as shown in Figure 1, the user at the desktop opportunistically transfers, via a Wi-Fi *ad hoc* link, a message for a friend to a user passing nearby, ‘hoping’ that this user will carry the information closer to the destination. This user passes close to a train station, and forwards the message to a traveller going to the same city where the destination user works. At the train station of the destination city, a car driver is going in the same neighbourhood of the destination’s working place. The driver meets the destination user on his way, and the message is finally delivered.

**Figure 1** The opportunistic networking concept (see online version for colours)



## 2.1 *Opportunistic networking case studies and applications*

Despite the fact that research on opportunistic networks dates back to just a few years ago, concrete applications and real case studies are already available (for a more extensive discussion about this point please refer to Pelusi, Passarella and Conti (2006a)).

The Huggle Project (<http://www.huggleproject.org>) is a four-year project, started in January 2006, funded by the European Commission in the framework of the FET-SAC initiative (<http://cordis.europa.eu/ist/fet/comms-sy.htm>). It targets solutions for

communication in autonomic opportunistic networks. Among the various activities, the project is putting emphasis on measuring and modelling pairwise contacts between devices. Pair-wise contacts between users and devices can be characterised by means of two parameters: contact durations and inter-contact times. The statistical properties of these parameters are used to drive the design of forwarding policies (Chaintreau et al., 2007). Furthermore, they are also at the basis of the design of concrete applications. For example, Huggle is working with epidemiologists to experimentally study the correlation between human contact patterns and the spread of diseases like flu. The patterns of contacts between people (measured in real experiments) are also the basis for designing ‘social-aware’ applications. An initial example of this approach is the design of a content distribution system in urban setting Leguay et al. (2006). Refined solutions for this type of applications are being designed in the Huggle project (e.g. Yoneki et al., 2007) thanks to the autonomic tools for detecting user social communities (Hui et al., 2007).

Opportunistic networks are also applied to interdisciplinary projects focusing on wildlife monitoring. Usually, small monitoring devices are attached to animals, and an opportunistic network is formed to gather information and carry it to a few Base Stations (BS) possibly connected to the Internet. Contacts among animals are exploited to aggregate data, and carry them closer and closer to the BS. This is a reliable, cost-effective and non-intrusive solution.

Concrete applications implementing these ideas have been used in the ZebraNet project (Juang et al., 2002). ZebraNet is an interdisciplinary project of the Princeton University performing novel studies of animal migrations and interspecies interactions, by deploying opportunistic networks on zebras in the vast savanna area of the central Kenya under control of the Mpala Research Centre (<http://www.mpala.org/researchctr/research/ongoing.html>).

Finally, we mention the use of opportunistic networks to bring Internet connectivity to rural areas. In developing countries and rural areas, deploying the infrastructure required to enable conventional Internet connectivity is typically not cost-effective. However, Internet connectivity is seen as one of the main booster to bridge the digital divide. Opportunistic networks represent an easy-to-deploy and extremely cheap solution. Typically, rural villages are equipped with a few collection points that temporarily store messages addressed to the Internet. Simple devices mounted on bus, bicycles or motorbikes that periodically pass by the village collect these messages and bring them in regions where conventional Internet connectivity is available (e.g. a nearby city), where they can be delivered through the internet. The same concept is exploited to enable communication in the opposite direction (from the Internet to villages). Projects implementing these concepts are currently ongoing. For example, the DakNet (Pentland, Fletcher and Hasson, 2004) and KioskNet (Guo et al., 2007) Projects focus on realising a very low-cost asynchronous ICT infrastructure to provide connectivity to rural villages in India, while the Saami Network Connectivity Project (Doria, Uden and Pandey, 2002) provides connectivity to inhabitants of Lapland.

### 3 Routing in opportunistic networks

In all the above case studies, routing is one of the most compelling challenges. The design of efficient routing strategies for opportunistic networks is generally a complicated task due to the absence of knowledge about the topological evolution of the network. Routing performance improves when more knowledge about the expected topology of the network can be exploited (Jain, Fall and Patra, 2004). Unfortunately, this kind of knowledge is not easily available, and a trade-off must be met between performance and knowledge requirement. A key piece of knowledge to design efficient routing protocols is information about the *context* in which the users communicate. Context information, such as the users' working address and institution, the probability of meeting with other users or visiting particular places, can be exploited to identify suitable forwarders based on context information about the destination. In the following of this sections, we classify the main routing approaches proposed in the literature based on the amount of knowledge about the context of users they exploit. Specifically, we identify three classes, corresponding to *context-oblivious*, *partially context-aware* and *fully context-aware* protocols.

#### 3.1 Context-oblivious routing

Basically, routing techniques in this class exploit some form of flooding. The heuristic behind this policy is that when there is knowledge neither of a possible path towards the destination nor of an appropriate next-hop node, a message should be disseminated as widely as possible. Protocols in this class might be the only solution when no context information is available. Clearly, they generate a high overhead (as we also highlight in the performance evaluation section), may suffer high contention and potentially lead to network congestion (Jindal and Psounis, 2007). To limit this overhead, the common technique is to control flooding by either limiting the number of copies allowed to exist in the network, or by limiting the maximum number of hops a message can travel. In the latter case, when no relaying is further allowed, a node can only send directly to the destination when (in case) it is met.

The most representative protocol of this type is Epidemic Routing (Epidemic for short; Vahdat and Becker, 2000). Whenever two nodes come into communication range, they exchange summary vectors that contain a compact unambiguous representation of the messages currently stored in the local buffers. Then, each node requests from the other the messages it is currently missing. The dissemination process is somehow bounded because each message is assigned a hop count limit giving the maximum number of hops it is allowed to traverse till the destination. When the hop count limit is set to one, the message can only be sent directly to the destination node.

Dissemination-based algorithms also include network-coding-based routing (Widmer and Le Boudec, 2005), which takes an original approach to limit message flooding. Just to give a classical example, let A, B and C, be the only three nodes of a string network, such as any message travelling between A and C has to be relayed by B. Let node A generate message  $a$  addressed to node C, and node C generate the message  $c$  addressed to node A. In a conventional forwarding scheme, node B has to relay message  $a$  to C and message  $c$  to A. In network coding, node B broadcasts a single packet containing  $a \oplus c$ . Once received  $a \oplus c$ , both nodes A and C can decode the messages. In

general, network coding-based routing outperforms flooding, as it is able to deliver the same amount of information with fewer messages injected into the network. A more extended survey about network coding techniques can be found in Pelusi, Passarella and Conti (2007).

An alternative, drastic way of reducing the overhead of Epidemic without relying on network coding is implemented by Spray and Wait (Spyropoulos, Psounis and Ragavendra, 2007). Message delivery is subdivided in two phases: the spray phase and the wait phase. During the spray phase, multiple copies of the same message are spread over the network both by the source node and those nodes that have first received the message from the source node itself. This phase ends when a given number of copies, say  $L$ , have been disseminated in the network. Then, in the wait phase each node holding a copy of the message (i.e. each relay node) stores its copy and eventually delivers it to the destination when (in case) it comes within reach. The analytical model derived in Spyropoulos, Psounis and Ragavendra (2007) shows that  $L$  can be chosen based on a target average delay. The spray phase may be performed in many ways. Under the assumption that nodes movements are independent and identically distributed, the *Binary Spray and Wait* policy is the best one in terms of delay. Any node (including the sender) holding  $n$  copies ( $n > 1$ ) of the message hands over  $\lfloor n/2 \rfloor$  copies to the first encountered node, and keeps the remaining copies for itself. When a node is left with only one copy of the message, it switches to direct transmission and only transmits the message to the final destination node when (if) it is met.

### 3.2 Partially context-aware routing

Partially, context-aware protocols exploit some particular piece of context information to optimise the forwarding task. The main difference with fully context-aware protocols is the fact that the latter usually provide a full-fledged set of algorithms to gather and manage any type of context information, while the former are customised for a specific type of context information.

Probabilistic Routing Protocol using History of Encounters and Transitivity (PROPHET; Lindgren, Doria and Schelen, 2003) is one of the most popular examples of protocols falling in this class. PROPHET is an evolution of Epidemic that introduces the concept of delivery predictability. The delivery predictability is the probability for a node to encounter a certain destination. The delivery predictability for a destination increases when the node meets the destination and decreases (according to an ageing function) between meetings. A transitivity law is also included in the algorithm, such that if node A frequently meets node B and node B frequently meets node C, then nodes A and C have high delivery predictability to each other. The PROPHET forwarding algorithm is similar to Epidemic except that, during a contact, nodes also exchange their delivery predictability to destinations of messages they store in their buffers, and messages are requested only if the delivery predictability of the requesting node is higher than that of the node currently storing the message.

The context information used by PROPHET is the frequency of meetings between nodes. The same type of context information is also used by MV (Burns, Brock and Levine, 2005) and MaxProp (Burgess et al., 2006), which, in addition, also exploit information about the frequency of visits to specific physical places. Other protocols use the time lag from the last meeting with a destination to estimate the probability of

delivering the messages. The bottomline idea (thoroughly investigated in Grossglauser and Vetterli (2003)) is that the decreasing gradient of the time lag identifies a suitable path towards the destination. Examples of protocols exploiting this piece of context information are Last Encounter routing (Grossglauser and Vetterli, 2003) and Spray and Focus (Spyropoulos, Psounis and Ragavendra, 2007).

In MobySpace routing (Leguay and Friedman, 2006), the mobility pattern of nodes is the context information used for routing. The protocol builds up a high dimensional Euclidean space, named MobySpace, where each axis represents a possible contact between a couple of nodes and the distance along an axis measures the probability of that contact to occur. Two nodes that have similar sets of contacts, and that experience those contacts with similar frequencies, are close in the MobySpace. The best forwarding node for a message is the node that is as close as possible to the destination node in this space. Obviously, in the virtual contact space just described, the knowledge of all the axes of the space also requires the knowledge of all the nodes that are circulating in the space. However, this full knowledge might not be required for successful routing.

The final example we mention is Bubble Rap (Hui and Crowcroft, 2007) in which the context information is the social community users belong to. In Bubble Rap, communities are automatically detected via the patterns of contacts between nodes. It is assumed that communities are labelled. Messages originating in a different community from the destination's one are forwarded as follows. Assume node A is carrying a message addressed to D, and meets node B. The message is handed over to B if the community of B is the same as the community of D, or if B has a higher ranking with respect to node A. The ranking is measured based on the set of peers a node is usually in touch with, and is thus a measure of the 'sociability' of nodes. Basically, Bubble Rap looks for nodes belonging to the same community of the destination. If such nodes are not found, it forwards the message to increasingly sociable nodes, which have more chances to get in touch with the community of the destination. Exploiting context information related to the social behaviour of people is one of the most promising research directions in the area.

### 3.3 Fully context-aware routing

Fully context-aware protocols not only exploit context information to optimise routing, but also provide general mechanisms to handle and use context information. The advantage of this approach is to be much more general than the approaches mentioned in Section 3.2. Indeed, these routing protocols can be used with *any* set of context information, thus allowing the system to be customised to the particular environment it has to operate in. To the best of our knowledge, two protocols only fall in this category, i.e. context-aware routing (Musolesi, Hailes, and Mascolo, 2005) and HiBOp (Boldrini, Conti and Passarella (2007a)).

context-aware routing assumes an underlying MANET routing protocol that connects together nodes in the same MANET cloud. To reach nodes outside the cloud, a sender looks for the node in its current cloud with the highest probability of delivering the message successfully to the destination. This node temporarily stores the message, waiting either to get in touch with the destination itself, or to enter a cloud with other nodes with higher probability of meeting the destination. Therefore, nodes in context-aware routing compute delivery probabilities proactively, and disseminate them in their *ad hoc* cloud. Note that context information is exploited to evaluate probabilities just for



those destinations each node is aware of (i.e. that happen to have been co-located in the same cloud at some time). The main focus of context-aware routing is on defining algorithms to combine context information (which is assumed available in some way) to compute delivery probabilities. Specifically, a multi-attribute utility-based framework is defined to this end. The framework is general enough to accommodate for different types of context information. As an example, in Musolesi, Hailes and Mascolo (2005) authors use residual battery life, the rate of connectivity change and the probability of meeting between nodes as context information.

With respect to context-aware routing, HiBOP is more general, as it does not necessarily require an underlying MANET routing protocol, and is able to exploit context information also for those nodes that have never been within the same cloud. Furthermore, the definition and management of context information is not addressed in context-aware routing, while it is a core part of HiBOP. Finally, and most importantly, context-aware routing does not capture, in the context definition, any information about the users social behaviour, which results in Boldrini, Conti and Passarella (2007a) demonstrate being a particularly valuable piece of information to design an efficient routing scheme.

Since, the performance analysis presented in this article focuses on the HiBOP protocol, we describe its mechanisms in more details in the following section.

### 3.4 The History-based Opportunistic Routing protocol

HiBOP is a fully context-aware routing protocol completely described in Boldrini, Conti and Passarella (2007a). HiBOP includes mechanisms to handle any type of context information. As a particular instance, in Boldrini, Conti and Passarella (2007a), the context is assumed to be a collection of information that describes the community in which the user lives, and the history of social 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 through nodes that share more and more context data with the message destination. Since users of HiBOP have possibly to share personal information, privacy issues should be considered. Privacy management in opportunistic networks is – in general – a topic still largely not addressed, and it is not the target of this article to provide complete privacy solutions for HiBOP. It should be noted that the set of information that is considered in Boldrini, Conti and Passarella (2007a) (and that we also consider hereafter) is equivalent to personal information people advertise on their public web pages (e.g. the working institution and address) which are, therefore, not perceived as sensitive information from a privacy standpoint. Designing complete privacy solutions for HiBOP is one of the main subjects of future work.

**Table 1** Identity Table

|        | <i>Personal information</i> |        | <i>Residence</i> |  |
|--------|-----------------------------|--------|------------------|--|
| Name   | John Doe                    | City   | Pisa             |  |
| E-mail | j.doe@iit.cnr.it            | Street | Via Garibaldi, 2 |  |
| ...    |                             | ...    |                  |  |

More in detail, HiBOP assumes that each node locally stores an Identity Table (IT), that contains personal information on the user that owns the device (an example is reported in Table 1). Nodes exchange ITs when getting in touch. At each node, its own IT, and the set of current neighbours' ITs, represent the *Current Context*, which provides a snapshot of the context the node is currently in.

The current context is useful in order to evaluate the *instantaneous* fitness of a node to be a forwarder. But even if a node is not a good forwarder because of its current location/neighbours, it could be a valid carrier because of its habits and past experiences. Under the assumption that humans are most of the time 'predictable', it is important to collect information about the context data seen by each node in the past, and the recurrence of these data in the node's current context. To this end, each context attribute seen in the current context (i.e. each row in neighbours' ITs) is recorded in a History Table, together with a continuity probability index, that represents the probability of encountering that attribute in the future (actually more indices are used, as described in Boldrini, Conti and Passarella (2007a)).

The main idea of HiBOP forwarding is looking for nodes that show increasing *match* with known context attributes of the destination. High match means high similarity between node's and destination's contexts and, therefore, high probability for the node to bring the message in the destination's community (possibly, to the destination). Therefore, a node wishing to send a message through HiBOP specifies (any subset of) the destination's IT in the message header. Any node in the path between the sender and the destination asks encountered nodes for their match with the destination attributes, and hands over the message if an encountered node shows a greater match than its own. The detailed algorithms to evaluate matches are described in Boldrini, Conti and Passarella (2007a). It is worth recalling here that matches are evaluated as delivery probabilities, and distinct probabilities are computed based on the current context ( $p_{CC}$ ) only, and on the History ( $p_H$ ) only. The final probability is evaluated via standard smoothed average, as  $p = \alpha \cdot p_H + (1 - \alpha) \cdot p_{CC}$ ,  $0 \leq \alpha \leq 1$ . The  $\alpha$  parameter allows HiBOP to tune the relative importance of the current context and History.

In HiBOP, just the source node is allowed to replicate the message, in order to tightly control the trade-off between reliability and message spread. Specifically, the source node replicates the message until the joint loss probability of nodes used for replication is below a system-defined threshold ( $p_l^{\max}$ ). Specifically, if  $p(i)$  is the delivery probability of the  $i$ th node used for replicating the message, and  $k$  is the number of nodes used for replication, the following equation holds:

$$k = \min \{j \mid \prod_{i=0}^j (1 - p_{(i)}) \leq p_l^{\max} \}.$$

#### 4 Performance of opportunistic routing approaches under social mobility patterns

The goal of this section is to compare the different opportunistic routing approaches in realistic human mobility scenarios. Specifically, we investigate the protocols' behaviour with respect to a number of parameters that describe user movement patterns. Therefore, before presenting the performance results, we briefly describe the mobility model we consider (Section 4.1). The performance evaluation is carried out by considering the

two opposite ends of the spectrum presented in Section 3. Specifically, we compare a context-oblivious routing protocol (Epidemic) with a fully context-aware routing protocol (HiBOP).

#### 4.1 Realistic mobility: the Home-cell Community-based Mobility Model

The HCMM (fully described in Boldrini et al. (2007)) is an evolution of the Community-based Mobility Model (CMM) proposed in Musolesi and Mascolo (2007). Community-based (or group) mobility models are attracting interest of researchers in the opportunistic networking area, because they are suitable to realistically model the influence of social relationships among people on the user mobility patterns.

As in CMM, in HCMM every node belongs to a social community (group). Nodes that are in the same social community are called *friends*, while nodes in different communities are called *non-friends*. Relationships between nodes are modelled through social links (each link has an associated weight). At the system start-up, all friends have a link to each other. Also two nodes that are not friends can have a link, according to the *rewiring probability* ( $p_r$ ) parameter. Specifically, for each node, each link towards a friend is rewired to a non-friend with  $p_r$  probability.

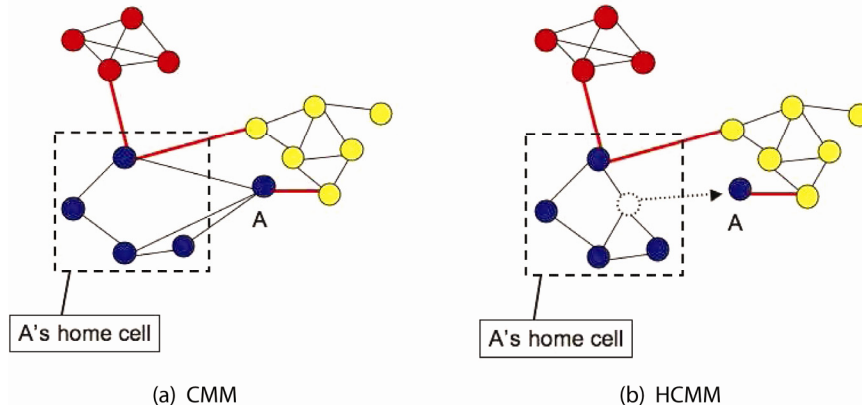
Social links are then used to drive node movements. Nodes move in a grid, and each community is initially randomly placed in a square of the grid. Nodes' movement is made up of two component: first, a node has to select the cell towards which to move. Node selects the target cell according to the social attraction exerted by each cell on the node. Attraction is measured as the sum of the links' weights between the node and the nodes currently moving in or towards the cell. Finally, the target cell is selected based on the probabilities defined by cells' attraction (i.e. if  $a_j$  is the attraction of cell  $j$ , then the probability of selecting that cell is  $a_j/\sum_j a_j$ ). After selecting the target cell, node selects the 'goal' within a cell (the precise point towards which node will be heading) according to a uniform distribution. Finally, speed is also selected accordingly to a uniform distribution within a user-specified range. HCMM (and CMM) also allows for collective group movements. Specifically, once every *reconfiguration period* nodes of each group select a (different) cell and move to that cell. Reconfigurations are synchronous across groups, i.e. all groups start moving to the new cell at the same time. Therefore, during reconfigurations nodes of different groups may get in touch.

The difference between HCMM and CMM is the way of considering the social relationships with nodes that are outside their starting cell (called 'home cell' in HCMM). Let us focus on Figure 2. In CMM, when node A moves outside its home cell, it 'carries over' all its social relationships, i.e. nodes that have social relationships with A are attracted towards the same cell towards which A is moving. In Boldrini et al. (2007), it is shown that this has an avalanche effect such that *all* nodes in A's home cell follow A.

This behaviour does not allow CMM to model relevant mobility patterns, because nodes are basically not only attracted by physical locations, but only by social relationships between each other. In HCMM, when A moves outside its home cell it does not carry over its social links. Nodes having social relationships with A are still attracted towards A's home cell. Furthermore, once A is outside its home cell, it selects its goal on next movements outside the home cell with probability  $p_e$ , and goes back to the home cell with probability  $1 - p_e$ . The rationale behind these modifications is the fact that there are several scenarios in which also physical locations (besides social relationships) play a

role in determining users' movements. In HCMM, people wishing to meet with A (i.e. having social attraction towards A) are attracted towards A's home cell because that is the most likely *physical* place where A can be met, or because their social relationship with A is conditioned to the fact that A is in its home cell (e.g. if someone want to meet an insurance agent, they will go the insurance office, not to the agent's house).

**Figure 2** Community-based Mobility Model vs. Home-cell Community-based Mobility Model (see online version for colours)



In a nutshell, HCMM models the fact that humans are social (belongs to groups), move towards other people they have relationships with (not only within their group, but also outside their group), and occasionally move collectively with their group. Furthermore, results presented in Boldrini et al. (2007) show that the duration of contact and inter-contact times under HCMM are similar to those measured in real experiments, which show that HCMM provides realistic movement traces.

#### 4.2 Performance evaluation strategy

In the following of the article, we highlight how the different routing approaches are able to autonomically react and adapt to the dynamically evolving conditions of the operating scenario. To this end, we exploit several control knobs provided by HCMM to highlight the different autonomic properties of Epidemic and HiBOP. Specifically, we identify three main reference cases for our study. In the first one (Section 4.3), we analyse the reactivity of routing protocols to sudden contacts among groups. Specifically, we focus on closed groups (i.e.  $p_r = 0$ ), and then we force groups to collectively move with varying frequency. Messages addressed to nodes outside the group can be delivered only during contacts between different group members during collective movements.<sup>2</sup> This analysis allows us to understand if routing protocols are able to exploit even those few chances to find good routes. We analyse this aspect by varying the reconfiguration interval parameter.

In the second scenario, (Section 4.4), we analyse the effect of social relationships between users. We want to understand how routing protocols react to different levels of users' sociality, measured as the probability of users having relationships outside their

reference group. Clearly, we achieve this by varying the rewiring parameter ( $p_r$ ). The higher  $p_r$ , the more nodes are ‘social’, the lesser groups are closed communities.

In the third scenario, we look at how protocols work in completely closed groups. In this case neither rewiring nor reconfigurations are allowed, and we place a different group in each cell of the grid. Therefore, the only chance of delivering messages between groups is by exploiting contacts between nodes at the borders of the cells. We study the routing protocols’ performance as function of the nodes’ transmission range. Basically, this scenario allows us to understand how protocols can exploit contacts that are not related to social relationships, but just happen because of physical co-location (e.g. contacts between people working for different companies in the same floor of a building).

We test routing performance in terms of *QoS perceived by users* and *resource consumption*. The user QoS is evaluated in terms of message delay and packet loss. Message delay is evaluated based on the first replica reaching the destination, while we count a packet loss if all replicas get lost. To highlight some specific different behaviour between Epidemic and HiBOP, in some cases we also show the average number of hops required by messages to be delivered, and we separate the delay for messages addressed to friend and non-friend nodes. Resource consumption is evaluated in terms of buffer occupation and bandwidth overhead. Specifically, the bandwidth overhead is computed as the ratio between the number of bytes generated in the whole network during a simulation run, and the number of bytes generated by the senders. Note that, we count in all overheads related to routing and forwarding, such as the exchanges of IT, requests for delivery probabilities, etc. To highlight specific differences, in a few cases we also show the number of copies spread in the network, and we separately highlight the bandwidth overhead related to data and non-data messages.

To highlight the effect of human mobility patterns only, we assume

- 1 infinite buffers
- 2 an ideal MAC level that completely avoids congestion impairments
- 3 an ideal physical channel where nodes experience 0% packet loss within a circular transmission range and 100% packet loss outside and
- 4 ‘infinite’ bandwidth (in the sense that messages can be always exchanged when nodes get in touch).

As thoroughly discussed in Boldrini, Conti and Passarella (2007a), this setup tends to favour dissemination-based schemes like Epidemic. More specifically, in this configuration HiBOP best results would be to approach the delay and packet loss achieved by Epidemic, while significantly reducing the resource consumption. Finally, unless otherwise stated, our setup consists of 30 nodes evenly divided in three groups. We assume a square simulation area  $1,250 \text{ m} \times 1,250 \text{ m}$  large, divided in a  $5 \times 5$  grid. The default transmission range is 125 m. Unless otherwise stated two nodes in each group generate messages, with an inter-generation time exponentially distributed (with average 300 sec). Each message is addressed to a friend or to a non-friend node with 50% probability. Messages expire after 18,000 sec. Each simulation run for 90,000 sec (of simulated time). For particular setups, we increased the run lengths so as to achieve a minimum amount of characteristic events in each run (e.g. reconfiguration runs with reconfiguration interval equal to 36,000 sec last for 397,000 sec). To make sure that


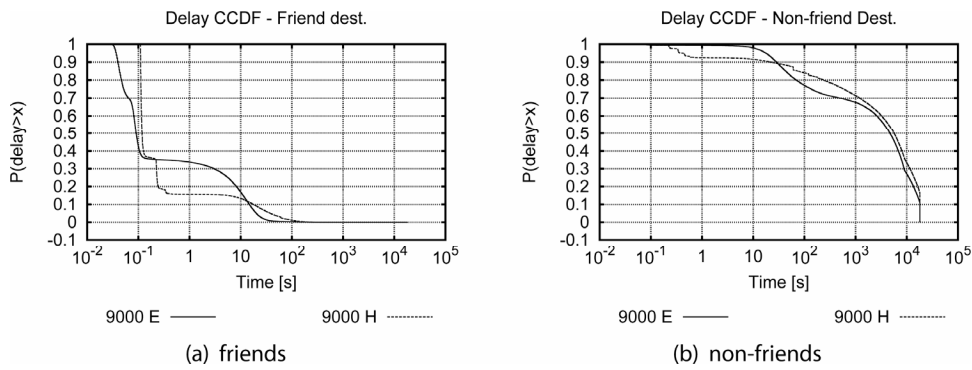
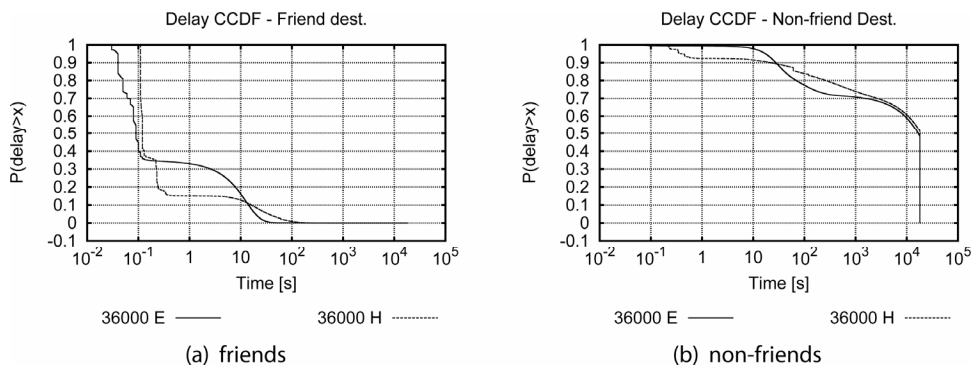
messages still not delivered at the end of a run will never be delivered (so as to achieve a correct measure of the packet loss index), during the last 18,000 sec, senders do not generate any new message. Furthermore, statistics are collected by eliminating the initial and final transitory regimes, i.e. by using the steady-state phase of simulation runs only. Each setup was replicated 50 times: statistics presented hereafter are averaged over the 50 replicas, with confidence interval at 95% confidence level.

### 4.3 *Impact of collective groups' movements (reconfigurations)*

It is worth recalling that in this scenario the rewiring probability is 0, and thus, except for reconfigurations, nodes do not have chances to meet. The reconfiguration interval varies among 2,250, 9,000 and 36,000 sec. Table 2 shows the QoS performance as a function of the reconfiguration interval. 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. The performance in terms of delay can be better highlighted by separately focusing on delay towards friend and non-friend nodes. Specifically, Figures 3–5 show the delay distribution towards friend nodes (left-hand-side plots) and non-friend nodes (right-hand-side plots) for the three reconfiguration periods. First of all, delays towards friends basically do not depend on the reconfiguration interval, since friends are always co-located in the same group. While only a small amount of messages destined to friend nodes experiences a delay greater than 10 sec, most (between 60 and 70% depending upon the reconfiguration interval) of the messages addressed to non-friend nodes experience a delay greater than 10<sup>3</sup> seconds. Furthermore, note that depending on the frequency of reconfigurations, distributions' tails are more or less 'heavy'. The worst case is clearly for a reconfiguration interval equal to 36,000 sec, where about 50% of messages towards non-friend destinations expire. Also note that, even though HiBOP provides higher packet loss and delay, the difference with Epidemic is quite thin. Note that, as buffers and bandwidth are not limited, Epidemic gives a reference upper bound on the performance achievable by any routing protocol. 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.

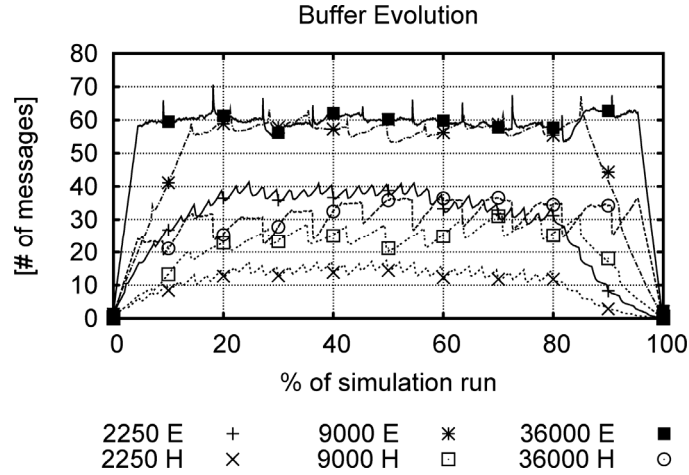
**Table 2** Users Quality of Service (focus on the reconfiguration parameter)

|                   | <i>Reconf (s)</i> | <i>HiBOP</i>     | <i>Epidemic</i>  |
|-------------------|-------------------|------------------|------------------|
|                   | 2,250             | 0 ± 0            | 0 ± 0            |
| <i>P</i> loss (%) | 9,000             | 8.16 ± 1.68      | 5.52 ± 1.46      |
|                   | 36,000            | 25.64 ± 1.30     | 24.12 ± 1.31     |
|                   | 2,250             | 1202.52 ± 91.09  | 907.10 ± 67.08   |
| Delay (s)         | 9,000             | 3651.68 ± 295.05 | 3204.58 ± 278.70 |
|                   | 36,000            | 5615.43 ± 225.93 | 5445.11 ± 161.53 |

**Figure 3** Delay distributions with reconfigurations every 2,250 sec**Figure 4** Delay distributions with reconfigurations every 9,000 sec**Figure 5** Delay distributions with reconfigurations every 36,000 sec

The good performance in terms of user QoS shown by HiBOp comes along with a drastic reduction in resource usage. Figure 6 shows the buffer occupation over time shown as a percentage of duration of a simulation run (points are average values over the replicas). HiBOp is much less greedy in spreading messages, and therefore the buffer occupation is drastically reduced. This is a general difference between Epidemic and HiBOp, which is confirmed in all scenarios we have tested. The extent of this reduction depends upon the scenario, and can be as high as an order of magnitude.

**Figure 6** Buffer occupation (focus on the reconfiguration parameter)



**Figure 7** Copies Distribution (focus on the reconfiguration parameter)


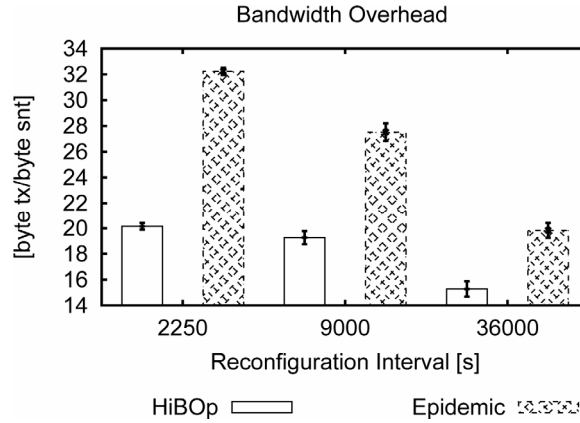


Figure 7 compares Epidemic and HiBOP with respect to the number of copies generated (recall that the number of nodes in the network is 30, thus the maximum number of copies is 29). High resource consumption for Epidemic is due to the fact that each node copies all its messages to all nodes it encounters. Therefore, the more the contacts between nodes, the more the spreading of messages. Figure 7 shows that approximately 50% of messages (corresponding to the messages with a non-friend destination) are spread by Epidemic across the *whole* network, when the reconfiguration interval is equal to 2,250 and 9,000 sec. The performance in terms of delay and packet loss shows that in this particular scenario flooding yields no significant advantages. As contacts during reconfigurations involve entire groups, a fully replication inside each group is not more convenient than replicating the message on a single node of each group. HiBOP, due to its reliability rule, tends to replicate the message inside the sender's group, but does not flood the other groups upon reconfigurations, thus resulting in lower number of copies.



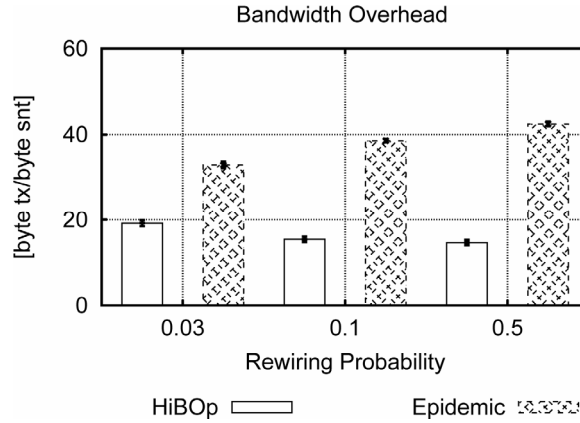
**Figure 8** Bandwidth overhead (focus on the reconfiguration parameter)

Finally, Figure 8 shows the bandwidth overhead of the two protocols. 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 36,000 sec down to 2,250 sec) 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. HiBOP behaves in a different way. When groups do not mix (reconfiguration interval equal to 36,000 sec) paths for messages going outside the sender’s group are seldom available. HiBOP realises this, because context information about nodes outside the group is rarely available, and avoids consuming resources uselessly. As nodes mix more and more (reconfiguration intervals equal to 9,000 and 2,250 sec), also HiBOP (as Epidemic) generates more overhead, 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 indicate that exploiting context information makes HiBOP much more efficient than flooding-based protocols, despite the additional resources needed for context management purposes.

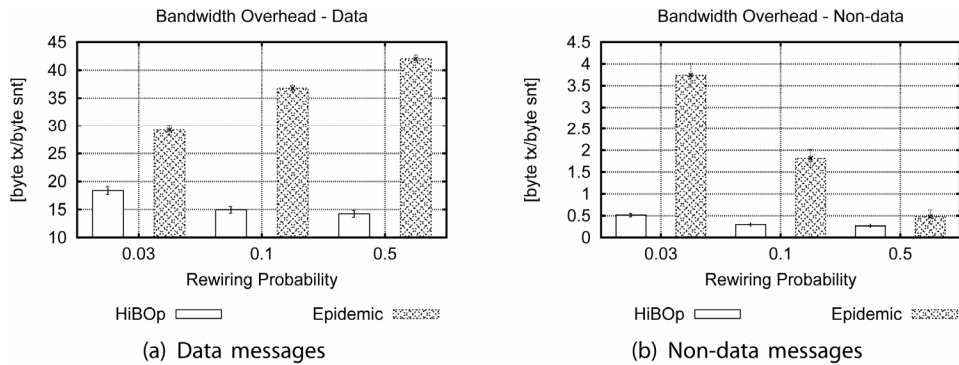
#### 4.4 Impact of user sociality

To understand the impact of the user sociality on routing performance, we vary the rewiring parameter ( $p_r$ ). When a node goes to a cell different from its home it shows to nodes in the ‘foreign’ cell context information related to its home cell, thus becoming a good next hop for messages destined to its friends. On the other hand, it roams in the foreign cell for a number of rounds and collects context data about nodes in that cell. When it then comes back to the home cell, this knowledge can effectively be used for sending messages to that particular foreign cell. Indeed, that node is likely to go back to the *same* foreign cell after a while, because the social links towards nodes in that cell are still active. Clearly, the routing performance is sensitive to the user sociality, because users having social relationships with other groups are the only possible way of getting messages out of the originating group. This sensitiveness impact differently on the

**Figure 9** Bandwidth overhead (focus on the rewiring parameter)



**Figure 10** Bandwidth overhead (focus on the rewiring parameter)



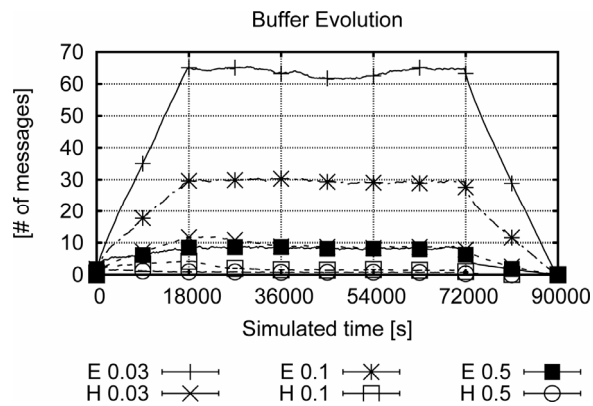
resource usage of HiBOP and Epidemic, as shown by Figure 9. Similar remarks drawn with respect to reconfiguration intervals apply also here. The higher the users sociality (high  $p_i$ ), the higher the mix between nodes and the forwarding opportunities. While Epidemic naively uses all these resources spreading messages, HiBOP leverages nodes' mixing (and the resulting spread of context information) to identify good paths more and more accurately.

Figure 10 shows how data and non-data traffic contributes to the bandwidth overhead. As already said, Epidemic exploits all the possibilities of reaching the destination by copying the messages on nodes as much as possible. This results in a high overhead, which is useless particularly for highly connected scenarios where there are a lot of forwarding opportunities. Note that the high Epidemic's overhead essentially comes from the aggressive replication of messages (i.e. from data traffic). Indeed, Figure 10(b) shows that the traffic related to forwarding (i.e. the traffic related to the exchange of summary vectors) actually decreases when more connectivity opportunities are available. The buffer occupation curves (Figure 11) indicate that for higher rewiring, the buffers under Epidemic are less full, because messages can be delivered more quickly to the destinations. Therefore, the size of summary vectors decreases, and this explains the


trend of Figure 10(b). However, the reduction in terms of forwarding traffic is overwhelmed by the aggressive spread of message, which results in an increase of the overhead related to the data traffic (Figure 10(a)) and, ultimately, to the overall overhead increase Figure 9. Unlike Epidemic, HiBOp ‘learns’ the degree of connectivity of the network and uses this knowledge for adjusting the load. More specifically, HiBOp learns the current state of the network through the exchange of context messages. As context information is spread more and more widely (rewiring equal to 0.1 and 0.5) paths become more and more known, and HiBOp reduces the exchanges of both data and non-data messages.

Epidemic’s high resources consumption is confirmed by Figure 12. With Epidemic, between 50 and 70% of messages are spread through the whole network. Epidemic tends to exploit all opportunities, regardless of the sociality of users. Therefore, when nodes are more mixed (higher rewiring), Epidemic floods the network more aggressively. As we will show when presenting the QoS performance figures, this is basically useless and thus results in wasting memory and bandwidth resources. HiBOp, instead, is aware of the current state of the network and adjusts the number of replicas of each packet based on the sociality of the network. Note that, even with the lowest sociality (rewiring = 0.03), only about 30% of messages are copied to more than 10 nodes. Note also that, unlike Epidemic, this percentage decreases to zero with higher levels of sociality.

**Figure 11** Buffer occupation (focus on the rewiring parameter)

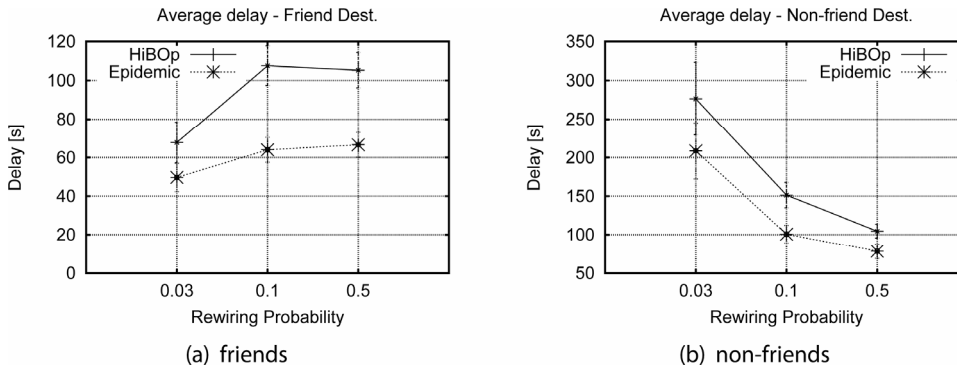
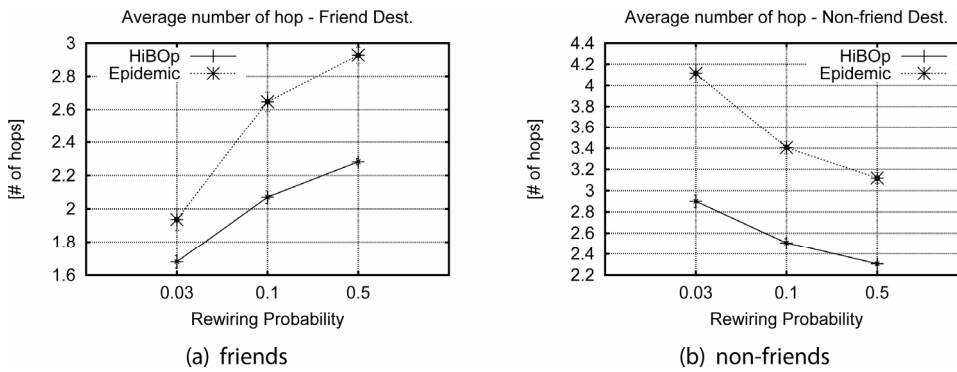


**Figure 12** Copies distribution (focus on the rewiring parameter)



**Table 3** Average delay (focus on the rewiring parameter)

|           | $p_r$ | HiBOP              | Epidemic           |
|-----------|-------|--------------------|--------------------|
| Delay (s) | 0.03  | $170.86 \pm 25.86$ | $130.28 \pm 20.59$ |
|           | 0.1   | $129.42 \pm 12.51$ | $83.20 \pm 8.57$   |
|           | 0.5   | $104.91 \pm 8.87$  | $73.69 \pm 7.16$   |

**Figure 13** Average delay (focus on the rewiring parameter)**Figure 14** Average number of hops (focus on the rewiring parameter)

As far as the QoS performance figures (Table 3), again the packet loss is negligible (so we do not show it), while – as expected – the average delay decreases as users become more social. The performance of HiBOP are still not far from the bound represented by Epidemic. It is also interesting to note (Figure 13) that the delay of messages towards friend nodes tends to slightly *increase* as users become more social, because they spend (on average) more time outside their home group. However, as shown by Table 3, the advantage of connecting more efficiently users between groups as users become more social overwhelms the slight performance reduction experienced by friends.

Mobility also affects the number of hops a message passes through before reaching its final destination (see Figure 14). As our setup simulates a *social* network, nodes belonging to the same community are expected to meet more frequently and for a longer time. This results in better QoS performances for messages destined to friends. As the network becomes more mixed, nodes tend to spend more time outside their community, thus becoming good forwarders for messages destined outside. The proximity between friends reduces as rewiring increases and more forwarding hops are needed in order to reach the destination (Figure 14(a)). On the other hand, the proximity between non-friend nodes increases and the number of hops a message passes through decreases (Figure 14(b)).

#### 4.5 Breaking closed groups

In this set of simulations, we use a  $3 \times 3$  grid with nine groups of five nodes each. Just one node, located in the upper left cell sends messages, destined to a node in the lower right cell. Recall that the only way a message can reach its final destination is through edge contacts with nodes between which no social relation exists. By varying nodes' transmission range we can analyse how this edge effect impacts on forwarding. We use three values for the transmission range, i.e. 62.5, 125 and 250 m. Therefore, nodes cover – on average – less than half a cell, slightly less than a cell, and one and a half cell.

The bottom line of the results is that HiBOP is not suitable for networks with no sociality. At very small transmission ranges (62.5 m), HiBOP is not able to deliver acceptable QoS (Table 4). HiBOP needs a minimum number of contacts between users to spread context information around. Indeed, at 125 m HiBOP restores acceptable QoS at least in terms of packet loss, and is fully effective at 250 m. Also in this case, Epidemic and HiBOP behave differently with respect to the bandwidth overhead (Figure 15). At 62.5 m, HiBOP seldom forwards messages. As context data is not circulating, nodes in the sender's group are almost all equally fit to carry the messages closer to the destination. At a high transmission range the context data is circulating effectively, and therefore good paths can be identified soon. In the intermediate cases (e.g. transmission range equal to 125 m), HiBOP is not (yet) able to correctly learn the status of the network, and this results in a higher overhead with respect to Epidemic. However, note that these results confirm that Epidemic is not able to exploit rich connectivity scenarios without flooding the network, since it increases its overhead at high transmission ranges.

**Table 4** Users Quality of Service (closed groups)

|                   | <i>Range (m)</i> | <i>HiBOP</i>       | <i>Epidemic</i> |
|-------------------|------------------|--------------------|-----------------|
| <i>P</i> loss (%) | 62.5             | 61.41 ± 10.16      | 0 ± 0           |
|                   | 125              | 0 ± 0              | 0 ± 0           |
|                   | 250              | 0 ± 0              | 0 ± 0           |
| Delay (s)         | 62.5             | 14732.57 ± 1242.74 | 535.50 ± 14.05  |
|                   | 125              | 576.40 ± 177.56    | 102.83 ± 1.82   |
|                   | 250              | 1.77 ± 0.55        | 23.58 ± 0.80    |

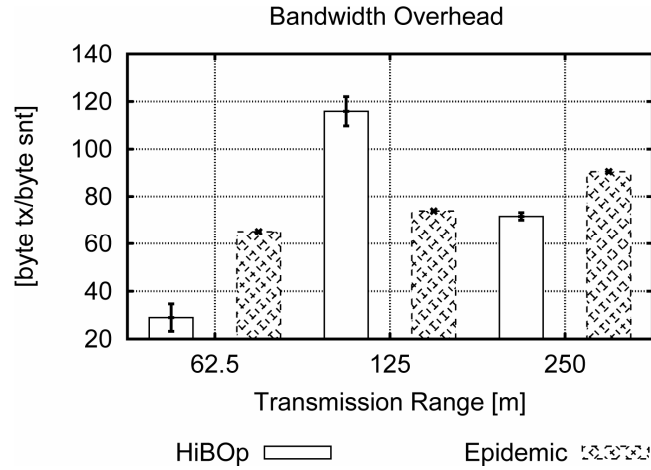

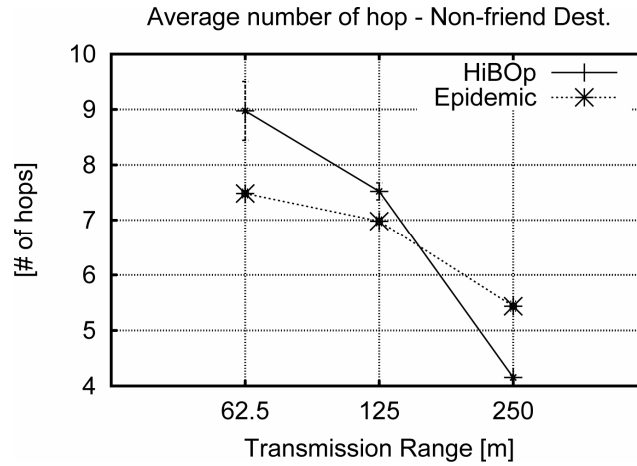
**Figure 15** Bandwidth overhead (closed groups)**Figure 16** Average number of copies (closed groups)

Figure 16 shows the average number of hops (recall that in this configuration statistics are related to non-friend nodes only). We can see that Epidemic generates 44 copies of each message, i.e. it replicates messages on all nodes, as it is not aware of the current state of the network. In HiBop, the number of copies increases as context information spreads, i.e. for increasing transmission ranges. This is because when the transmission range is low there is no reason to replicate messages, since no good paths can be found in a context-aware scheme if context information cannot spread. As soon as context information can be exploited, paths can be found and HiBop starts replicating messages. Finally, Figure 17 shows the average number of hops. In both cases, this figure decreases with higher transmission ranges, as more contact opportunities become available, and a single hop is able to bring messages closer to the destination.

**Figure 17** Average number of hops (closed groups)

## 5 Conclusions

In this article, we have highlighted how different approaches to routing in opportunistic networks are able to autonomously adapt to the dynamic scenarios resulting from humans' mobility patterns. We have framed this work in the ongoing research on routing for opportunistic networks, and we have compared the performance figures of two protocols at the opposite ends of the spectrum as far as the use of context information, namely Epidemic and HiBOp.

Our main findings can be summarised as follows. Context-based routing actually provides an effective congestion control mechanism, and, with respect to dissemination-based routing, provides acceptable QoS with drastically lower overhead, unless in very adverse scenarios. Indeed, HiBOp is able to automatically learn the connectivity opportunities determined by users movement patterns, and exploit them efficiently. This autonomic, self-learning feature is completely absent in dissemination-based routing schemes.

Our results also suggest a hybrid scheme for networks with varying levels of user sociality. When groups are very isolated, context data cannot circulate, and cannot be used for taking effective forwarding decisions. In such cases, dissemination-based schemes seem the only way to enable communication between groups. As soon as users become more social, context information spreads in the network, and 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 select the appropriate routing scheme.

From a complementary standpoint, our results show that in opportunistic networks *user sociality helps routing*: users' relationships outside their 'home' community allow context information to spread in the network, and make forwarding more and more efficient. These results open interesting research directions. Actually, since opportunistic networks build the network by exploiting mobile devices people carry with them, looking at social network theories to model users' social relationships and exploit these models

for designing network protocols is a very interesting research direction. Indeed, the EC FET-PERADA SOCIALNETS project (due to start in February 2008) will be looking at these aspects. Other interesting research directions include providing privacy and security support through distributed and scalable systems in opportunistic networks. Also, another challenging research direction is how to integrate purely infrastructure-less opportunistic networks (like the ones we have considered in this article) with access points to the internet infrastructure. Finally, designing data-management systems (built on top of opportunistic routing schemes) to improve data availability in opportunistic networks is another direction we find extremely important.

### Acknowledgements

This work was partially funded by the European Commission under the HAGGLE (027918) FET-SAC Project, and under the SOCIALNETS (217141) FET-PERADA Project.

### References

- Boldrini, C., Conti, M. and Passarella, A. (2007a) ‘Impact of social mobility on routing protocols for opportunistic networks’, Paper presented in the Proceedings of the *First IEEE WoWMoM Workshop on Autonomic and Opportunistic Networking (AOC 2007)*.
- Boldrini, C., Conti, M. and Passarella, A. (2007b) ‘Users mobility models for opportunistic networks: the role of physical locations’, Paper presented in the Proceedings of the *IEEE Wireless Rural and Emergency Communications Conference (WRECOM 2007)*.
- Boldrini, C., Conti, M., Iacopini, I. and Passarella, A. (2007) ‘Hi-BOp: a history based routing protocol for opportunistic networks’, Paper presented in the Proceedings of the *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2007)*, pp.1–12.
- Borgia, E., Conti, M., Delmastro, F. and Pelusi, L. (2005) ‘Lessons from an *Ad-Hoc* network test-bed: middleware and routing issues’, *Ad Hoc and Sensor Wireless Networks, an International Journal*, Vol. 1.
- Burgess, J., Gallagher, B., Jensen, D. and Levine, B. (2006) ‘MaxProp: routing for vehicle-based disruption-tolerant networks’, Paper presented in the Proceedings of the *25th IEEE Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2006)*.
- Burns, B., Brock, O. and Levine, B. (2005) ‘MV routing and capacity building in disruption tolerant networks’, Paper presented in the Proceedings of the *24th IEEE Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2005)*.
- Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R. and Scott, J. (2007) ‘Impact of human mobility on opportunistic forwarding algorithms’, *IEEE Transactions on Mobile Computing*, Vol. 6, pp.606–620.
- Conti, M. and Giordano, S. (2007a) ‘Multihop *Ad Hoc* networking: the reality’, *IEEE Communications Magazine*, Vol. 45, pp.88–95.
- Conti, M. and Giordano, S. (2007b) ‘Multihop *Ad Hoc* networking: the theory’, *IEEE Communications Magazine*, Vol. 45, pp.78–86.
- Doria, A., Uden, M. and Pandey, D. (2002) ‘Providing connectivity to the Saami nomadic community’, Paper presented in the Proceedings of the *2nd International Conference on Open Collaborative Design for Sustainable Innovation (DYD 2002)*.
- Fall, K. (2003) ‘A delay-tolerant network architecture for challenged internets’, Paper presented in the Proceedings of the *2003 ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM 2003)*, pp.27–34.



- Grossglauser, M. and Vetterli, M. (2003) 'Locating nodes with EASE: last encounter routing in ad hoc networks through mobility diffusion', Paper presented in the Proceedings of the *22nd IEEE Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE INFOCOM 2003)*.
- Guo, S., Falaki, M.H., Oliver, E.A., Rahman, S.U., Seth, A., Zaharia, M.A. and Keshav, S. (2007) 'Very low-cost internet access using kiosnet', *SIGCOMM Computer Communication Review*, Vol. 37, pp.95–100.
- Hui, P. and Crowcroft, J. (2007) 'Bubble rap: forwarding in small world dtns in every decreasing circles', *Technical report, Technical Report UCAM-CL-TR684*. Cambridge, UK: University of Cambridge.
- Hui, P., Yoneki, E., Chan, S. and Crowcroft, J. (2007) 'Distributed community detection in delay tolerant networks', Paper presented in the Proceedings of the *2nd ACM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch 2007)*.
- Jain, S., Fall, K. and Patra, R. (2004) 'Routing in a delay tolerant network', Paper presented in the Proceedings of the *2004 ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM 2004)*, pp.145–158, New York, NY: ACM.
- Jindal, A. and Psounis, K. (2007) 'Contention-aware analysis of routing schemes for mobile opportunistic networks', Paper presented in the Proceedings of the *1st International ACM MobiSys Workshop on Mobile Opportunistic Networking (MobiOpp 2007)*, pp.1–8.
- Juang, P., Oki, H., Wang, Y., Martonosi, M., Peh, L. and Rubenstein, D. (2002) 'Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with ZebraNet', *ACM SIGPLAN Notices*, Vol. 37, pp.96–107.
- Leguay, J. and Friedman, T.C.V. (2006) 'Evaluating mobility pattern space routing for dtns', Paper presented in the Proceedings of the *25th IEEE Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2006)*, pp.1–10.
- Leguay, J., Lindgren, A., Scott, J., Friedman, T. and Crowcroft, J. (2006) 'Opportunistic content distribution in an urban setting', Paper presented in the Proceedings of the *2006 SIGCOMM Workshop on Challenged Networks (CHANTS 2006)*, pp.205–212.
- Lindgren, A., Doria, A. and Schelen, O. (2003) 'Probabilistic routing in intermittently connected networks', *ACM Mobile Computing and Communications Review*, Vol. 7, pp.19–20.
- Musolesi, M. and Mascolo, C. (2007) 'Designing mobility models based on social network theory', *ACM SIGMOBILE Mobile Computing and Communication Review*, Vol. 11, pp.59–70.
- Musolesi, M., Hailes, S. and Mascolo, C. (2005) 'Adaptive routing for intermittently connected mobile ad hoc networks', Paper presented in the Proceedings of the *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2005)*, pp.183–189.
- Pelusi, L., Passarella, A. and Conti, M. (2006a) 'Beyond manets: dissertation on opportunistic networking', *Technical Report, IIT-CNR*, Available at: [http://bruno1.iit.cnr.it/~andrea/tr/commag06\\_tr.pdf](http://bruno1.iit.cnr.it/~andrea/tr/commag06_tr.pdf).
- Pelusi, L., Passarella, A. and Conti, M. (2006b) 'Opportunistic networking: data forwarding in disconnected mobile ad hoc networks', *IEEE Communications Magazine*, Vol. 44, pp.134–141.
- Pelusi, L., Passarella, A. and Conti, M. (2007) *Handbook of Wireless Ad Hoc and Sensor Networks, Chapter Encoding for Efficient Data Distribution in Multi-Hop Ad Hoc Networks*. New York, NY: Wiley and Sons Publisher.
- Pentland, A., Fletcher, R. and Hasson, A. (2004) 'DakNet: rethinking connectivity in developing nations', *IEEE Computer*, Vol. 37, pp.78–83.
- Spyropoulos, T., Psounis, K. and Ragavendra, C. (2007) 'Efficient routing in intermittently connected mobile networks: the multiple-copy case', *ACM/IEEE Transactions on Networking*, Vol. 16, pp.77–90.
- Vahdat, A. and Becker, D. (2000) *Epidemic Routing for Partially Connected Ad Hoc Networks. Technical Report CS-2000-06*, Computer Science Department. Duke University.

- Widmer, J. and Le Boudec, J. (2005) 'Network coding for efficient communication in extreme networks', Paper presented in the Proceedings of the *ACM SIGCOMM 2005 Workshop on Delay Tolerant Networking (WDTN 2005)*, pp.284–291.
- Yoneki, E., Hui, P., Chan, S. and Crowcroft, J. (2007) 'A socio-aware overlay for publish/subscribe communication in delay tolerant networks', Paper presented in the Proceedings of the *10th ACM Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems (MSWiM 2007)*, pp.225–234.
- Zhang, Z. (2006) 'Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: overview and challenges', *IEEE Communications Surveys and Tutorials*, Vol. 8, pp.24–37.

## Notes

<sup>1</sup>As will be clear in the following, in opportunistic networks the routing and forwarding tasks are strictly intertwined and usually performed at the same time. Therefore, hereafter we use the terms routing and forwarding interchangeably.

<sup>2</sup>The probability of contacts due to groups choosing adjacent cells is typically low due to the high number of cells with respect to the number of groups.