The Stability Region of the Delay in Pareto Opportunistic Networks

Chiara Boldrini, Marco Conti, *Member, IEEE*, and Andrea Passarella, *Member, IEEE*

Abstract—The intermeeting time, i.e., the time between two consecutive contacts between a pair of nodes, plays a fundamental role in the delay of messages in opportunistic networks. A desirable property of message delay is that its expectation is finite, so that the performance of the system can be predicted. Unfortunately, when intermeeting times feature a Pareto distribution, this property does not always hold. In this paper, assuming heterogeneous mobility and Pareto intermeeting times, we provide a detailed analysis of the conditions for the expectation of message delay to be finite (i.e., to converge) when social-oblivious or social-aware forwarding schemes are used. More specifically, we consider different classes of social-oblivious and social-aware schemes, based on the number of hops allowed and the number of copies generated. Our main finding is that, in terms of convergence, allowing more than two hops may provide advantages only in the social-aware case. At the same time, we show that using a multi-copy scheme can in general improve the convergence of the expected delay. We also compare social-oblivious and social-aware strategies from the convergence standpoint and we prove that, depending on the mobility scenario considered, social-aware schemes may achieve convergence while social-oblivious cannot, and vice versa. Finally, we apply the derived convergence conditions to three popular contact data sets available in the literature (Cambridge, Infocom, and RollerNet), assessing the convergence of each class of forwarding protocols in these three cases.

Index Terms—Opportunistic networks, DTN, routing protocols, performance models, delay convergence

✦

1 INTRODUCTION

T HE great popularity of the delay tolerant networking paradigm is due to its ability to cope with challenged network conditions, such as high node mobility, variable connectivity, and disconnected subnetworks, that would impair communications in traditional mobile ad hoc networks. Opportunistic networks are an instance of the delay tolerant paradigm applied to networks made up of users' portable devices (such as smartphones and tablets). In this scenario, user mobility becomes one of the main drivers to enable message delivery. In fact, according to the store-carry-and-forward paradigm, user devices store messages and carry them around while they move in the network, exchanging them upon encounter with other nodes, and eventually delivering them to their destination.

An opportunistic forwarding protocol defines the strategy according to which messages are exchanged during encounters. Two main approaches can be identified: *social-oblivious* protocols, which do not exploit any information about the users' context and social behaviour, but just hand over the message to the first node encountered, and *social-aware*[1] protocols, which make use of information on how users behave or which social relations they share in order to make predictions on users' future behavior that might be useful for forwarding messages. Depending on the number of copies generated for the same message, forwarding protocols can be further classified into single-copy or multi-copy schemes. Forwarding protocols may also differ in the number of intermediate hops that they exploit. Simpler strategies may be single-hop or two-hop strategies (e.g., Direct Transmission and Two Hop [2]), while others can allow multi-hop paths to bring the message to the destination.

Modelling the performance of social-oblivious and social-aware forwarding protocols for opportunistic networks is still an open research issue. Knowing the distribution of intermeeting times and the rules applied by the forwarding algorithm used in the network, one could—in principle—model the distribution of the delay experienced by messages and compute its expectation. In practice, modeling analytically the delay of the various forwarding protocols for general distributions of inter meeting times is very hard, and models exist only for some specific cases, typically assuming exponential intermeeting times [1], [3], [4], [5], [6], [7]. A related modeling challenge is to assess the convergence of routing protocols, i.e., whether a specific protocol yields finite or infinite expected delay. Assessing convergence allows us to understand whether a particular protocol can be safely used or not given a pattern of intermeeting times and how to configure it so that it converges, if possible. Although less informative than a complete delay model, convergence models can be derived for a large class of routing protocols releasing the exponential intermeeting time assumption, as shown in this paper.

The convergence of the expected delay is not guaranteed in all cases in which the expectation of the intermeeting times may diverge. In fact, being the delay the result of the composition of the time intervals between node encounters, depending on the convergence of intermeeting times,

---

1. These policies are also referred to as utility-based [1], in contrast to randomized strategies, corresponding to our social-oblivious schemes.

• *The authors are with the Institute for Informatics and Telematics of the Italian National Research Council, Via G. Moruzzi 1, 56124 Pisa, Italy. E-mail: {chiara.boldrini, marco.conti, andrea.passarella}@iit.cnr.it.*

the expectation of the delay itself might diverge. This can happen, for example, when intermeeting times feature a Pareto (also known as power law) distribution, as first highlighted in [8]. The problem with Pareto distributions is that their expectation is finite only for certain values of their exponent α. More specifically, the expectation is finite if $\alpha > 1$, while for $\alpha \leq 1$ it diverges to infinity. The first to postulate the existence of Pareto intermeeting times in real mobility scenarios (i.e., analyzing real traces of human mobility) were Chaintreau et al. in their seminal work in [8]. The relevance of Pareto intermeeting times in opportunistic networks is both theoretical and empirical. Cai and Eun [9] have mathematically derived that heavy-tailed intermeeting times can emerge depending on the relationship between the size of the boundary of the considered scenario and the relevant timescale of the network, showing that, at least in principle, Pareto intermeeting times are something that one may be faced with when studying opportunistic networks. Empirical evidence for the presence of Pareto intermeeting times was first suggested by [8], but it has been later criticised, arguing that the tail of the distribution is in fact exponential (e.g., [10]). Typically, these results are derived focusing on the aggregate intercontact time distribution, while convergence depends on pairwise distributions. As proved in [11], the aggregate and pairwise distributions can be in general very different, and therefore analysis of pairwise inter-contact times are necessary, which are however mostly missing in the literature. To address this issue, we have performed a pairwise hypothesis testing on three popular publicly available contact data sets (Cambridge, Infocom'05, and RollerNet—see Section 8) and we have found that the Pareto hypothesis for intermeeting times cannot be rejected for 80, 97, and 85.5 percent of pairs, respectively. We believe that these results provide a strong case for Pareto intermeeting times in opportunistic networks and substantially motivate analyses like the one presented in this paper.

Under the Pareto intermeeting times assumption, in this work we derive the stability region (i.e., the Pareto exponent values of pairwise intermeeting times for which finite expected delay is achieved) of a broad class of social-oblivious and social-aware forwarding protocols (single- and multi-copy, single- and multi-hop). The starting point of our paper is the work by Chaintreau et al. [8], where such conditions have been studied for the two-hop scheme (see Section 2 for more details) under the assumption of homogeneous mobility (i.e., i.i.d. intermeeting times across all pairs). However, measurement studies [12], [8] have shown that real networks are intrinsically heterogeneous. Thus, in this paper, we investigate whether heterogeneity in contact patterns helps the convergence of the expected delay of a general class of social-oblivious (Section 5) and social-aware (Section 6) forwarding protocols, and whether convergence conditions can be improved using multi-copy strategies and/or multi-hop paths. In general, we find that there is no protocol or family of protocols that always outperform the others (Section 7). More specifically, the key findings presented in the paper are the following:

- For *social-oblivious strategies*, if convergence can be achieved, *two hops are enough* for achieving it.

- Using n *hops* can help *social-aware schemes*, and make them converge in some cases when all other social-aware or social-oblivious schemes diverge.

- In both the social-oblivious and the social-aware case, we find that *multi-copy strategies* can achieve a finite expected delay even when single-copy strategies cannot.

- Comparing *social-oblivious and social-aware multi-copy solutions*, we are able to prove mathematically that there is no clear winner between the two, since either one can achieve convergence when the other one fails, depending on the underlying mobility scenario.

In addition to these results, we discuss the related work in Section 2 and the network model we refer to in Section 3. Our reference forwarding policies are described in Section 4. Finally, in Section 8 we showcase the main results of this work by applying the derived convergence conditions to three popular contact data sets available in the literature.

2 RELATED WORK

This work is orthogonal to the literature on models of delay in opportunistic networks, since we provide the conditions for the *existence* of a finite delay. Once convergence has been verified, the expected delay value can be computed using complete delay models like the one in [13]. End-to-end delay models are typically much more difficult to obtain than convergence models, and therefore convergence models are a very useful first step in the analytical characterisation of forwarding protocols in opportunistic networks. Most existing models assume that intermeeting times are approximately exponentially distributed [5], [6], [14], and in these cases convergence is never an issue. However, when this assumption does not hold, convergence becomes a critical evaluation aspect, and should be studied preliminarily to any additional analysis of the exact value of the expected delay. To the best of our knowledge, there is no other contribution, besides that of Chaintreau et al. [8], that considers the problem of the convergence of the expected delay when intermeeting times feature a Pareto distribution. Our work differs from that of Chaintreau et al. both in the mobility settings and in the forwarding schemes considered. More specifically, we focus on the more general case of heterogeneous intermeeting times (as opposed to the homogeneous mobility considered in [8]), we extend the set of social-oblivious policies considered and we add the social-aware case.

Forwarding protocols for opportunistic networks can be classified as social-oblivious or social-aware protocols, depending on whether they use information on the way nodes behave in order to make forwarding decisions. In this paper we abstract the detailed mechanisms of both classes of protocols, in order to study their convergence properties, as discussed in Section 4. The simplest social-oblivious protocol is Direct Transmission [2], in which the source node is only allowed to deliver the message directly to the destination, if ever encountered. At the opposite side of the spectrum, with Epidemic routing [15] a new copy of the message is generated and handed over (both by the source and intermediate relays) any time a new node is encountered. In order to mitigate the side effects of Epidemic-style forwarding schemes in resource constrained environments, controlled flooding

solutions have been proposed (e.g., Spray&Wait [3], gossiping [7]). Another popular social-oblivious forwarding protocol is the Two Hop scheme [2], in which a message is forwarded by the source node to the first node encountered, which is then allowed only to pass the message directly to the destination. The Two Hop strategy has been shown to guarantee the maximum capacity in a homogeneous network [2].

Social-aware strategies can have different levels of awareness. Simplest approaches exploit information such as time since the last encounter (Spray&Focus [3]) or frequency of encounters (PROPHET [16]). This information is used to predict future meetings between pairs of nodes and thus to select relays that can guarantee a quick delivery according to the heuristic in use. In more complex strategies, the centrality of nodes in the social graph connecting the users of the network is used as an indicator of the ability to deliver messages (see, e.g., BUBBLE [17], SimBet [18]). Alternatively, as in the case of HiBOp [19] and SocialCast [20], the fitness of a node as a forwarder is computed from information on the context the users live in, e.g., information on the people they meet, the friends they have, the places they visit.

This paper extends our previous work in [21], which was only focused on social-oblivious forwarding strategies. In this work, besides extending the convergence conditions for the m-copy 2-hop case that we derived in [21], we include the analysis of social-aware forwarding strategies and a detailed comparison between social-aware and social-oblivious strategies from the convergence standpoint.

3   PRELIMINARIES

Our model considers a network with N mobile nodes. For the sake of simplicity, we hereafter assume that messages can be exchanged only at the beginning of a contact between a pair of nodes and that the transmission of the relayed messages can be always completed within the duration of a contact. In addition, we assume that each message is a bundle [22], an atomic unit that cannot be fragmented. We also assume infinite buffer space on nodes. All the above assumptions allow us to isolate, and thus focus on, the effects of node mobility from other effects, and are common assumptions in the literature on opportunistic networks modelling (they are used in most of the literature reviewed in Section 2). In addition, for the sake of comparison with [8], we also assume that the probability that two nodes meet is greater than zero for all node pairs. This ensures that, in principle, all nodes can meet with each other. Therefore, cases of deadlock (a message reaches a node which is impossible to leave due to the total absence of contacts with either other possible relays or the destination) are not possible. The only cause of divergent expected delay are the distributions of intermeeting times.

As we neglect transmission time, the actual duration of the contact is not critical. Thus, the main role in the experienced delay is played by intermeeting times, which are defined as the time between two consecutive meetings between the same pair of nodes.[2] We denote with M_{ij} the

intermeeting times between nodes i and j. For the sake of tractability, we assume that the network is stationary and that intermeeting times for a specific node pair i, j are i.i.d. . Under these assumptions, the encounter process between two nodes i and j can be seen as a renewal process with renewal intervals distributed as M_{ij}.

The message generation process and the mobility process are assumed to be independent. Thus, the time at which a new message is generated can be treated as a random time in the evolution of the mobility process, and thus the message sees the network as an observer arriving at a random point in time would. For this reason, in our analysis we will often use the concept of residual intermeeting time, defined as the time two nodes that are not in contact at a random time t_0 have to wait before meeting again. We denote the residual intermeeting time for the i, j node pair as R_{ij}.

Under our assumption of Pareto intermeeting times, the intermeeting time M_{ij} between a generic pair of nodes i and j is described by the CCDF $F_{M_{ij}}(t) = (\frac{t_{min_{ij}}}{t + t_{min_{ij}}})^{\alpha_{ij}}$, in which we use the definition of the Pareto distribution that allows for values arbitrarily close to zero, usually denoted as American Pareto [23] or Pareto distribution of the second kind [24].[3] Parameters $\alpha_{ij}(>0)$ and $t_{min_{ij}}$ are the shape and scale of the Pareto distribution and, similarly to the reference literature [8], [10], in the following we restrict to the case of power law random variables having the same scale, i.e., $t_{min_{ij}} = t_{min}, \forall i, j$. We will use the following properties of Pareto intermeeting times throughout the paper (please refer to [25, Appendix A, available in the online supplementary material] for more details):

P1 $E[M_{ij}]$ converges (i.e., is finite) if and only if (iif) $\alpha_{ij} > 1$.

P2   The residual intermeeting times R_{ij} associated to M_{ij} feature an American Pareto distribution with shape $\alpha_{ij} - 1$ and scale t_{min} [23], hence their expectation converges if and only if $\alpha_{ij} > 2$.

P3 $\min_j\{R_{ij}\} \sim \text{Pareto}(\sum_j(\alpha_{ij} - 1), t_{min})$ and its expectation converges if and only if $\sum_j(\alpha_{ij} - 1) > 1$.

P4   Assume $R_{ij}^{t_i}$ denotes the residual conditioned to the fact that i and j met at time t_i. Then the expectation of $R_{ij}^{t_i}$ converges if and only if the expectation of R_{ij} converges.

Furthermore, in the mathematical analysis in Sections 5 and 6, we will also heavily rely on the result in Lemma 1 below. The intuition behind it is the following. In the general case, the time before a node i currently holding a copy of the message hands it over to another node j depends on whether nodes i and j met in time interval (t_0, t_i), where t_0 is the message generation time. In fact, meetings correspond to renewals in the encounter renewal process between i and j, hence, from the meeting time on, we

---

2. Without loss of generality, here we assume a deterministic unit disk graph model for radio propagation, which is a common assumption in the literature on opportunistic networks. The proposed framework still applies for every other model of radio propagation.

3. The stability region derived in this paper holds also for the other version of the Pareto distribution, usually denoted as European Pareto, as discussed in [25]. Content in [25] not included in this paper is provided as supplemental material, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TMC.2014.2316506.

should consider the intermeeting time and not its residual. However, Lemma 1 below tells us that, when intermeeting times feature a Pareto distribution, we can simply study the case in which nodes i and j did not meet in (t_0, t_i) (i.e., model the time to the next encounter as a residual time), thus simplifying our analysis (for more details, see the proof of [25], Lemma 1]).

Lemma 1 (Worst-case waiting time). *Assume that node i has received a copy of the message at time t_i. In the worst case (happening with a non negligible probability), the time before node i hands over the message to another node j can be modeled as $R_{ij}^{t_i - t_0}$ (i.e., R_{ij} conditioned to be greater than $t_i - t_0$) or, equivalently from a convergence standpoint, as R_{ij}.*

4 FORWARDING STRATEGIES

In this section we summarise the main variants of opportunistic forwarding schemes that will be later evaluated against each other as far as the convergence of their expected delay is concerned. We identify two main strategies that forwarding protocols can adopt in order to improve their forwarding performance, namely the number of copies generated and the number of hops allowed. As we show later on in the section, it is easy to place any of the most popular routing protocols proposed in the literature in this classification.

As far as multi-copy strategies are concerned, here we only allow the source node to create and hand over multiple copies. Other possible configurations (e.g., intermediate relays allowed to generate new copies, like in the Spray&Wait case [3]) are left as future work. With respect to the number of hops n, we assume it to be either limited arbitrarily (e.g., using the TTL field) or naturally constrained by the forwarding strategy (e.g., social-aware schemes can exploit a number of intermediate relays that is at most equal to the number of nodes that are better forwarders—according to some social-aware metric—than the source node). In all cases, the last relay can only deliver the message to the destination directly. Please note that in this paper we only consider the case in which both the source node and intermediate relays refuse the custody of copies that they have already relayed (i.e., we assume that nodes are *memoryful*).[4] For this to be feasible, we assume that the identity of previous relays is enclosed into the copy's header. In the case of multiple copies, we assume that the source node does not use the same relays multiple times, and that relays do not accept the custody of the same copy of the message more than once. They can be used, however, as relays for different copies of the same message (as avoiding this would need to keep track of all forwarded messages at each relay, which would make protocols not scalable).

Due to the variety of social-aware schemes available in the literature (see Section 2) and the limited space, here we only consider an abstract social-aware protocol that measures how good a relay is for a given destination in terms of

4. In [21] we have derived the convergence conditions for the memoryless version of the class of social-oblivious forwarding protocols considered here, showing that the absence of memory always penalizes the convergence.

its *fitness*. The fitness fit_i^d is assumed to be a function of how often node i meets the destination d, thus fit_i^d can be taken as proportional to the rate of encounter $\frac{1}{E[M_{id}]}$ between node i and the destination. Under this abstract and general social-aware strategy, upon encounter, a node i can hand over the message to another node j only if its fitness is lower than the fitness of the peer, i.e., if $fit_j^d > fit_i^d$ holds (in the following we drop superscript d). The fitness function considered here uses only information on contacts between nodes, which have a direct dependence on the intermeeting time distribution. This lets us clearly show what is the impact of the contact dynamics on the performance of opportunistic forwarding protocols. How such simple fitness function can be extended to more complex forwarding strategies has been discussed in [13].

The combinations of the forwarding characteristics described above can be found in well known routing strategies. For example, the 1-hop 1-copy forwarding corresponds to Direct Transmission [2], while the 2-hop 1-copy forwarding is equivalent to the Two Hop forwarding introduced in [2]. The 2-hop m-copy forwarding is equivalent to the multi-copy version of the Two Hop protocol studied in [8]. Note that for most of the social-aware protocols, the number of copies and the maximum hops are also defined as parameters of the algorithm.

5 EXPECTED DELAY CONVERGENCE FOR SOCIAL-OBLIVIOUS SCHEMES

In this section we study under which conditions the expected delay of the social-oblivious schemes described in Section 4 converges for a *tagged* source-destination pair. We denote it as $E[D_{sd}]$, where s and d are the source and destination node, respectively. Simultaneous convergence for all source-destination pairs simply requires combining the conditions derived in the paper.

Recall that according to social-oblivious forwarding a message is handed over to the first feasible relay encountered. In the following, we denote with $\mathcal{P}_i$ the set of all nodes that can be encountered by node i. For the sake of comparison with [8], we assume that the probability of an encounter between any pair of nodes is strictly greater than zero (hence, we have that $|\mathcal{P}_i| = N - 1$ for all nodes i) and that $\alpha_{ij} > 1$ for all i, j node pairs.

5.1 Single-Copy Schemes

Theorem 1 below focuses on the 1-copy 1-hop social-oblivious scheme, which corresponds to the popular Direct Transmission scheme. In the following, we omit the proof since this result follows directly from property P2 and we move to the analysis of the 1-copy 2-hop scheme immediately.

Theorem 1 (1-copy 1-hop scheme). *$E[D_{sd}]$ converges if and only if $\alpha_{sd} > 2$.*

Theorem 2 (1-copy 2-hop scheme). *$E[D_{sd}]$ converges if and only if both the following conditions hold true:*

C1 $\sum_{j \in \mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$,

C2 $\alpha_{jd} > 2, \forall j \in \mathcal{P}_s - \{d\}$.

Proof. The protocol converges if and only if both the delay at the first hop converges and the delay at the second hop

converges. We analyse the former, first. The delay of the first hop converges iff the time required by the source to hand over the message, which is the time to encounter the first node in set $\mathcal{P}_s$, is finite. The source node s can either deliver the message directly to the destination or hand it over to an intermediate relay. The time before the source node hands over the message is distributed as $\min_{j \in \mathcal{P}_s}\{R_{sj}\}$, which is the time before the first node (possibly including the destination) is encountered. From property P3, we know that $\min_{j \in \mathcal{P}_s}\{R_{sj}\}$ has a finite expectation iff $\sum_{j \in \mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$, thus obtaining condition C1. We now consider the convergence of the second hop. As social oblivious protocols cannot control which relay is used, the delay of the second hop is finite iff delays from all possible relays to the destinations are finite. If the node to which the message has been handed over is not the destination but another generic node j, the expected delay from j to d is finite iff the expectation of the time before j meets d is finite. Exploiting Lemma 1, we can model the time before node j hands over the message to d as R_{jd}, whose expectation is finite iff $\alpha_{jd} > 2$. Please note that from here on, due to lack of space, we will not prove again the necessity of the convergence conditions we derive. In all cases it will be straightforward to prove it using the same argument outlined above. The complete proofs are however available in [25]. □

According to Theorem 1, the Direct Transmission protocol yields a convergent expected delay only if the source node meets the destination with a residual intermeeting time whose expectation converges. This clearly follows from the fact that the source node cannot exploit any other relays for the forwarding of the message. In the case of the two-hop scheme, the expectation converges even if the source node is not able to ensure convergence with a direct delivery. This can happen if the source node is able to hand over the message to any of the possible relays within a convergent expected time (Condition C1) and if the meeting process between this relay and the destination has a residual whose expectation converges (Condition C2). Please note that condition C1 alleviates the convergence condition on the source node at the expense of the additional condition C2 on intermediate relays.

With Theorem 3 we extend the analysis of single-copy schemes by studying their n-hop version.

Theorem 3 (1-copy n-hop scheme). *$E[D_{sd}]$ converges if and only if conditions C1 and C2 in Theorem 2 hold true.*

Proof. See of [25, Appendix C, available in the online supplementary material] for a complete proof. The intuitive reason behind Theorem 3 is that, since the first hop (from source to first relay) and last hop (from last relay to destination) are equivalent to those in Theorem 2, they also share the same convergence conditions (C1 and C2). For intermediate hops, it is possible to prove that convergence conditions are looser than C1 and C2, which are then sufficient and necessary for convergence. □

Theorem 3 tells us that, when using single-copy social-oblivious schemes, letting the message traverse more than two hops does not improve the convergence of the expected delay. Thus, when convergence is the only goal, network resources can be saved using the two-hop social-oblivious scheme without impairing the convergence of the expected delay.

5.2 Multi-Copy Schemes

As discussed in Section 2, when multiple copies of the same message can travel in parallel the opportunities to reach the destination are multiplied. In this section we investigate whether this also positively affects the convergence of the expected delay. Please note that hereafter we discuss the complete proof only for the first lemma, which provides the rationale for obtaining the other results of this section (for which just an intuitive proof is sketched, while more details can be found in [25, Appendix C, available in the online supplementary material]).

5.2.1 Two-Hop Forwarding

Recall that, according to the multi-copy version of the two-hop forwarding scheme, the source node hands over a copy of the message to the first m encountered nodes, which will then be only allowed to deliver the message directly to the destination, if ever encountered. In Lemmas 2 and 3 we study separately the first hop and the second hop, then putting together their results in Theorem 4. The goal is to derive how many convergent copies the source node can send out at the first hop and how many are needed for having a convergent second hop. In fact, as we demonstrate below, the higher the number of copies on the intermediate relays, the easier the convergence at the second hop. Thus, the number of copies that the source node is able to hand over within a finite expected time is critical to the convergence of the whole path. It is possible to prove that first-hop convergence becomes more difficult as the number of available relays decreases. Hence, after a certain point, the number of relays left does not allow to deliver one more copy within a finite expected time, setting an upper bound on the maximum number of first-hop convergent copies that the source node can send. This number depends also on the order in which relays are used (i.e., on the Pareto exponents of the available relays), which in turn depends on the sequence of encounters at the source node. Clearly, this order cannot be controlled and it is only the result of the evolution of the meeting process. Since the source node can meet at most $N - 1$ nodes, the possible sequences of distinct encounters are $(N - 1)!$. Let us denote as π_i the ith of these permutations. For each possible permutation π_i, in Lemma 2 we are able to compute the maximum number (max_i^{so}) of convergent copies that can be sent at the first hop by the social-oblivious source node. Then, considering all possible permutations π_i, we can identify (Corollary 1) a range of values (specifically, $[max_{lo}^{so}, max_{up}^{so}]$) within which max_i^{so} can vary, and under which permutations π_i the extreme values of the interval are achieved.

Lemma 2 (max_i^{so}). *When intermediate relays are selected by the source node according to sequence π_i, the source node is able to deliver at most max_i^{so} copies to as many relays with finite first hop expected delay, with max_i^{so} being equal to the following:*

$$max_i^{so} = \underset{m}{\arg\min}\left\{f_{max}^{so}(m, \pi_i) > 0\right\}, \tag{1}$$

where $f_{max}^{so}(m, \pi_i) = m + \sum_{z=m}^{|\mathcal{P}_s|} \alpha_z^{(i)} - (2 + |\mathcal{P}_s|)$ and $\alpha_z^{(i)}$ denotes the α_{sj} exponent of the zth node belonging to π_i.

Proof. At the first hop m copies are relayed to the first m distinct encountered nodes. Thus, the delivery process at the first hop is a selection without repetitions: every time a relay is selected, it is removed from the set of future relays for the same message.

Let us define $\mathcal{P}_s^k$ as the set of relays still available to s when the source node is delivering the kth copy, t_0 the time at which the message is generated at the source, and t_k the time at which the kth copy is handed over. Given that we assume that the probability that any two nodes meet is greater than zero, we have that $|\mathcal{P}_s| = N - 1$ and $|\mathcal{P}_s^k| = N - 1 - (k-1) = N - k$. Exploiting Lemma 1, the time before the kth copy is relayed is given by $\min_{j \in \mathcal{P}_s^k}\{R_{sj}\}$, which converges (Property P3) as long as $\sum_{j \in \mathcal{P}_s^k}(\alpha_{sj} - 1) > 1$, or equivalently, $\sum_{j \in \mathcal{P}_s^k} \alpha_{sj} > 1 + |\mathcal{P}_s^k|$, with $|\mathcal{P}_s^k| = N - k$. In order to achieve convergence for the m copies, this condition should be satisfied for all k from 1 to m.

We start by finding whether convergence is achieved for a fixed m. Lemma C1 in [25] tells us that the smaller the cardinality of the set of random variables of which we take the minimum, the tougher the convergence. This implies that the strictest condition for the convergence of the expected delay of the first hop is imposed by the mth copy, i.e., by the one that sees the smallest set of nodes left for relaying. Thus, if we are able to define a convergence condition for the mth copy, then it follows that the finiteness of the expected time to relaying for all previous copies is automatically guaranteed. Let us thus focus on the relaying of the mth copy. When the $(m-1)$th copy has been delivered, there are $N - 1 - (m - 1) = N - m$ potential relays left for the mth copy. The identities of these $N - m$ potential relays depend on the previous evolution of the forwarding process (i.e., which nodes have already been used). More specifically, there can be $(N-1)!$ different[5] permutations of the $N-1$ nodes in $\mathcal{P}_s$, while there can be $\binom{N-1}{N-m}$ possible combinations for the relays in $\mathcal{P}_s^m$. Let us denote with π_i the ith of the $(N-1)!$ permutations and with ν_i its corresponding combination. That is, taken sequence $\pi_i = \{a, e, c, b, f, h, g\}$ of encounters (where a, b, c, e, f, g, h are the nodes that the source node can meet) and assuming $m = 3$ we denote with ν_i the set $\{c, b, f, h, g\}$, i.e, the set of nodes available as relays once the first and second copies have been handed over. Let us now define a mapping $g^{(i)}$ that goes from set $\{\alpha_{sj}\}_{j \in \mathcal{P}_s}$ to set $\{\alpha_z^{(i)}\}_{z \in \{1, \ldots, |\mathcal{P}_s|\}}$, where $\alpha_z^{(i)}$ corresponds to the exponent α_{sj} of the zth element in π_i. Using the above notation, the time before the mth copy is handed over is described by $\min_{j \in \nu_i} R_{sj}$. Using again property P3 and the mapping defined above, we have

that the convergence condition for the expected delay of the mth copy is given by the following:

$$\sum_{z=m}^{|\mathcal{P}_s|} \alpha_z^{(i)} + m - (N + 1) > 0. \tag{2}$$

As discussed before, since the mth copy experiences the worst conditions for convergence, guaranteeing convergence for the mth copy implies automatic convergence of all previous copies. Hence, Equation (2) characterizes the stability region for first hop convergence.

The above equation defines the convergence condition for the mth copy when relays are encountered according to encounter sequence π_i. For a given node permutation π_i, we can also compute the greatest m value for which convergence is achieved, and in the following we discuss how. Recall that, according to [25, Lemma C1, available in the online supplementary material] convergence becomes more difficult as m increases. This is highlighted also by Equation (2). In fact, the left-hand side of the equation (hereafter denoted as $f_{max}^{so}(m, \pi_i)$) decreases as m increases (the formal demonstration is at the end of the proof). This implies that either $f_{max}^{so}(m, \pi_i)$ is always above/below zero or $f_{max}^{so}(m, \pi_i)$ crosses the x-axis at a certain point. If $f_{max}^{so}(m, \pi_i)$ is always below zero, the source node is not able to send any copy with finite first hop expected delay. Otherwise, the maximum number of convergent copies (for a given node encounter sequence π_i) that the source node can send is equal to the greatest m for which $f_{max}^{so}(m, \pi_i)$ is still above zero. Hence, Equation (1) follows.

To conclude the proof, let us now demonstrate that $f_{max}^{so}(m, \pi_i)$ decreases with m. To this aim, consider moving from m to $m + 1$. Function $f_{max}^{so}(m + 1, \pi_i)$ can be rewritten as $\sum_{z=m}^{|\mathcal{P}_s|} \alpha_z^{(i)} - \alpha_m^{(i)} + m + 1 - (N + 1)$. Thus, the difference between $f_{max}^{so}(m + 1, \pi_i)$ and $f_{max}^{so}(m, \pi_i)$ is $1 - \alpha_m^{(i)}$. $1 - \alpha_m^{(i)}$ is always smaller than zero, since we have assumed $\alpha_{ij} > 1$ for all i, j node pairs. This implies that the left-hand side of Equation (2) decreases as m increases. □

Corollary 1. *Quantity max_i^{so} derived in Lemma 2 takes values in the interval $[max_{lo}^{so}, max_{up}^{so}]$. The upper and lower bound on max_i^{so} (corresponding to the best and worst case for convergence) are reached when π_i corresponds to nodes encountered in increasing and decreasing order of α_{sj}, respectively.*

Proof. Let us provide an intuitive explanation for this result. We can divide the set $\mathcal{P}_s$ of possible relays at the source node into two disjoint sets, one containing the nodes that have already been used as relays and one containing those that have not. Clearly, as copies are handed over by the source node, nodes move from the second subset to the first subset. Convergence is determined by the exponents associated with nodes in the second subset (nodes still to be encountered). The higher the exponents in this subset, the easier the convergence, and vice versa. When convergence is easier, the source node can send more copies. Conversely, when convergence is more difficult less convergent copies can be sent. Thus, in the best case the exponents associated with nodes in the second

---

5. Please note that any of these permutations happen with non negligible probability, since we assume that all nodes can meet with each other. A rigorous proof can be obtained exploiting the same argument used in Part 4 of the proof of Theorem 3.

subset are the highest among the nodes in $\mathcal{P}_s$, while in the worst case such exponents are the lowest. From this, Corollary 1 follows.   □

Let us now focus on the second hop. The sequence π_i according to which the source node meets the other nodes affects not only the first hop delay but also the second hop delay. In fact, the relays picked by the source node according to π_i are those that are in charge of bringing the message to its final destination. It is possible to prove [25, Lemma C1, available in the online supplementary material] that the higher the number of relays the easier the convergence. However, given a sequence of encounters π_i, there exists a minimum number of relays that is enough for guaranteeing convergence. We denote this number as min_i^{so}, and we derive it in Lemma 3.

Lemma 3 (min_i^{so}). *Assuming that intermediate relays are selected in the order specified by sequence π_i, the expected delay from intermediate relays to the destination d will converge if and only if there are at least min_i^{so} intermediate relays, with min_i^{so} being equal to the following:*

$$min_i^{so} = \arg\min_m \{f_{min}^{so}(m, \pi_i) > 0\}, \qquad (3)$$

where $f_{min}^{so}(m, \pi_i)$ is defined as $\sum_{z=1}^m \alpha_z^{(i)} - (1 + m)$ and $\alpha_z^{(i)}$ denotes the exponent α_{jd} associated with the zth node in encounter sequence $\pi_i - \{d\}$.

Proof. The second hop can be modelled as a parallel delivery from m relays to the destination. Let us consider the ith relay, assuming that it receives its copy of the message at time t_i. The time before the ith relay hands over its copy to the destination can be modeled as a residual intermeeting time (Lemma 1). Considering all m relays, the time before the first of the m copies reaches the destination can be modeled as the minimum of the residual intermeeting times between the relays and the destination. Once we focus on a specific sequence $\pi_i - \{d\}$ of encounters at the source node, it is clear that the first m relays correspond to the first m nodes in the sequence. We denote with $\alpha_z^{(i)}$ the α_{jd} exponent associated with the zth node in π_i. Then, applying property P3, we obtain the convergence condition $\sum_{z=1}^m \alpha_z^{(i)} - m > 1$. Since, as discussed before, convergence becomes easier as m increases, the minimum number min_i^{so} of copies required at the second hop for convergence under sequence of encounters $\pi_i - \{d\}$ corresponds to the first (integer) m value in the above equation for which the condition is satisfied. Hence Equation (3) follows. For a detailed proof, see [25].   □

Corollary 2. *Quantity min_i^{so} derived in Lemma 3 takes values in $[min_{lo}^{so}, min_{up}^{so}]$. The upper and lower bounds on min_i^{so} (corresponding to the worst and best case for convergence) are reached when π_i corresponds to the sequence of nodes ordered in increasing and decreasing order of their exponents α_{jd}, respectively.*

Now, we use the results in Lemmas 2 and 3 for deriving the stability region of the delay under m-copy 2-hop social-oblivious forwarding (with $m < N - 1$).

Theorem 4 (m-copy 2-hop scheme). *$E[D_{sd}]$ converges if and only if the following condition holds true:*

C3 $m \geq min_{up}^{so} \wedge max_i^{so} \geq min_i^{so}, \forall i \in \{1, \dots, |\mathcal{P}_s^p|\}$,

where set $\mathcal{P}_s^p$ is the set of all permutations for elements in $\mathcal{P}_s$.

Proof. In order to derive C3, first we notice that, e.g., the first hop and second hop worst case (max_{lo}^{so} and min_{up}^{so}, respectively) in general do not happen simultaneously. In fact, meeting processes between nodes are independent, and the fact that node j meets the destination frequently (high α_{jd}) does not generally imply that it also meets the source node frequently (high α_{sj}), and vice versa. Thus, the order on set $\{\alpha_{sj}\}$ determined by sequence π_i does not correspond to the same ordering on set $\{\alpha_{jd}\}$. Since worst cases are not correlated (either positively or negatively), we have to impose convergence on all the possible combinations for relay selections. This implies deriving a sequence of α_{sj} based on π_i and its corresponding sequence of α_{jd} and verifying convergence for each of these permutations. In practice, we compute a pair (max_i^{so}, min_i^{so}) for each possible sequence π_i of relays. Convergence is possible as long as $max_i^{so} \geq min_i^{so}$ for all permutations, since this means that the first hop is always able to provide to the second hop the number of copies needed for convergence. When the above condition is satisfied, convergence is ensured as long as we send a number m of copies equal to or greater than the number of copies needed in the worst case at the second hop, hence[6] we set $m \geq min_{up}^{so}$.   □

Corollary 3. *A sufficient condition for the convergence of the expected delay under the memoryful m-copy two-hop forwarding scheme in Theorem 4 is given by the following:*

$$C3_{[s]} \quad m \geq min_{up}^{so} \wedge min_{up}^{so} \leq max_{lo}^{so}.$$

Proof. The sufficient condition $C3_{[s]}$ follows directly from Lemmas 2 and 3. What these lemmas told us is that, in the worst case, the first hop delivery can at most provide max_{lo}^{so} copies (with finite first hop expected delay) while, again in the worst case, the second hop delivery needs at least min_{up}^{so} copies. When $C3_{[s]}$ holds true, it is guaranteed that, in all cases, the minimum number of copies needed at the second hop is provided by the first hop, thus proving the sufficiency of the condition.   □

As discussed before, Chaintreau et al. [8] studied the m-copy two-hop scheme under homogeneous mobility patterns (corresponding to $\alpha_{ij} = \alpha, \forall i, j$). For the sake of completeness, in [25, (Appendix C.2.1), available in the online supplementary material] we verify that Theorem 4 confirms and extends the results in [8].

5.2.2   Multi-Hop Forwarding

Again we consider a social-oblivious protocol in which the source node generates m copies of the message and hands them over to the first m nodes encountered. Once the source

---

6. Please note that m can be configured to be smaller or greater than max_i^{so} for a given π_i. In the first case, the source node will simply send m convergent copies rather than max_i^{so}. In the second case, the source node will be able to send max_i^{so} with finite first hop expected delay and all other copies will be divergent.

TABLE 1
Summary of Convergence Conditions for Social-Oblivious and Social-Aware Routing Strategies

	1 hop		2 hops		n-hop	
	1 copy	m copies	1 copy	m copies	1 copy	m copies
social-oblivious	$\alpha_{sd} > 2$	-	[C1,C2]	[C3]	[C1,C2]	[C1,C2]
social-aware	$\alpha_{sd} > 2$	-	[C4,C5]	[C8]	[C6, C7]	[C6, C7]

node has handed over the m copies, these copies travel along multi-hop social-oblivious paths until the destination is found. Theorem 5 describes the convergence conditions that apply in this case.

Theorem 5 (m-copy n-hop scheme). *$E[D_{sd}]$ converges if and only if condition C1 and C2 in Theorem 2 hold true.*

Proof. As we did before, we only sketch the proof and we refer the reader to [25, Appendix C, available in the online supplementary material] for the rigorous mathematical derivation. Here, the source node is memoryful and thus it guarantees that the m copies are relayed to m distinct nodes. However, it is possible to prove that, after the first hop, there is a non negligible probability that all m copies are relayed to the same node. This is clearly a worst case as far as the convergence of the expected delay is concerned, because the parallel delivery offered by the multi-copy approach is not exploited. Since basically the multi-copy forwarding process turns into a 1-copy n-hop scheme, it means that copies in addition to the first one are useless in terms of convergence. Thus, we simply need to ensure that at least one copy achieves convergence, which is guaranteed by the same conditions applying to the 1-copy n-hop scheme, i.e., C1 and C2. □

5.3 Discussion

Table 1 summarises the results derived so far for social-oblivious forwarding protocols. In the following we will informally speak about *convergent* relays to denote nodes for which the associated convergence condition is satisfied. The first interesting finding is that n-hop social-oblivious protocols (last two columns of Table 1) are no more effective in delivering the message with finite expected delay than the simple 1-copy 2-hop forwarding. In fact, they both share the same convergence conditions (C1 and C2), but the former consumes much more network resources than the latter. So, if we are only interested in the convergence of the expected delay, paths with more than two hops should be avoided, as two hops ensure that the available forwarding diversity between nodes is explored, while minimizing resource consumption.

With social-oblivious protocols, when the source node meets the destination with a residual intermeeting time having $\alpha_{sd} > 2$, there is no reason to exploit other relays, as this will only introduce the chance of picking a bad relay. In fact, when the number of hops is allowed to grow, we have to impose on intermediate relays additional constraints that are not needed by Direct Transmission (see, e.g., condition C2 in Theorem 3 which requires that the residual intermeeting time between any relay and the destination achieves a finite expectation).

Different is the situation in which $\alpha_{sd} \leq 2$. In this case, the source node cannot reach destination d directly with a finite

expected delay but it may be able to hand over the message to other nodes within a finite expected time. Hence, exploring more relays can prove convenient. If these intermediate relays are *all* able to individually deliver the message to the destination within a finite expected time, then the 1-copy 2-hop strategy guarantees convergence while minimizing resource consumption.

Instead, when there exists at least one divergent intermediate relay, the most effective strategy is the m-copy 2-hop forwarding, under which the source is able to send up to max_{up}^{so} copies of the message. If $max_{up}^{so} = 1$, we find again conditions C1 and C2 that hold for the 1-copy 2-hop strategy. However, if the source node can reach operating point $max_{up}^{so} > 1$, conditions on the delivery from the relays to the destination become less restrictive since the more the copies sent out by the destination (with finite first hop expected delay) the easier the convergence at the second hop (Lemma 3).

6 EXPECTED DELAY CONVERGENCE FOR SOCIAL-AWARE SCHEMES

In this section our goal is to derive the convergence conditions for the social-aware approaches introduced in Section 4, which will then be used to investigate whether the social-aware approach outperforms the best social-oblivious ones. In the following, we denote with $\mathcal{R}_i$ the set of possible relays for node i, i.e., the set of nodes whose fitness is greater than that of node i. Recall that, with social-aware forwarding, nodes can hand over a message only to nodes with higher fitness.

6.1 Single-Copy Schemes

We start our discussion with the case of single copy schemes. Please recall that social-aware strategies do not make sense when only one hop is allowed, since this hop is necessarily the destination itself and Theorem 1 holds. Thus we go straight to the 1-copy 2-hop case.

Theorem 6 (1-copy 2-hop social-aware scheme). *$E[D_{sd}]$ converges if and only if the following conditions hold:*

C4 $\sum_{j \in \mathcal{R}_s} \alpha_{sj} > 1 + |\mathcal{R}_s|$,
C5 $\alpha_{jd} > 2, \forall j \in \mathcal{R}_s - \{d\}$.

Proof. The proof is a step-by-step repetition of the proof of Theorem 2, with the only difference that this time relays belong to $\mathcal{R}_s$, thus we omit the proof. □

Theorem 6 mirrors Theorem 2 with the exception that only nodes with fitness higher than that of the source node can be selected. At first sight, this seems only a minor difference, but it proves extremely significant in all those cases in which the source node is already a "good" relay (from the convergence standpoint). In these cases, in fact,

with social-aware forwarding we are sure that only relays better than the source node can be picked, thus ensuring that convergence can only improve, never get worse, at the second hop.

Theorem 7 (1-copy n-hop social-aware scheme). $E[D_{sd}]$ converges if and only if the following condition holds:

C6. $\sum_{j \in \mathcal{R}_i} \alpha_{ij} > 1 + |\mathcal{R}_i|$ for all $i \in \mathcal{R}_s \cup \{s\}$,
C7. $n \geq |\mathcal{D}| + 1$,

where set $\mathcal{D}$ comprises nodes $j \in \mathcal{R}_s$ whose exponent value α_{jd} is smaller than or equal to 2.

Proof. The proof exploits the ordering guaranteed by social-aware policies. Specifically, when social-aware policies are used, messages are forwarded along a path with increasing fitness. For the sake of simplicity, in the following we assume that there cannot be two nodes with the same fitness value. Recalling that $\mathcal{R}_i$ denotes the set of potential relays when the message is on node i, we have that, for a generic path $\{s, i, \ldots, j, z, d\}$ with increasing fitness, the relation $\mathcal{R}_s \supset \mathcal{R}_i \supset \cdots \mathcal{R}_j \supset \mathcal{R}_z$ holds (Fig. 1). Exploiting Lemma 1, we know that the time before the message leaves a generic node i is distributed as $\min_{j \in \mathcal{R}_i}\{R_{ij}\}$. According to property P3, the above expression has a finite expectation if and only if $\sum_{j \in \mathcal{R}_i} \alpha_{ij} > 1 + |\mathcal{R}_i|$ (condition C6).

In order to complete the proof, we have to consider the fact that when a message has reached the maximum number $n - 1$ of allowed intermediate hops, the relay currently holding the message can only deliver it to the destination directly. Thus, $\alpha_{jd} > 2$ is required after $n - 1$ relays have been reached. Let us split all possible relays in $\mathcal{R}_s$ into two subsets $\mathcal{C}$ and $\mathcal{D}$, such that $\mathcal{C} \cup \mathcal{D} = \mathcal{R}_s$. Subset $\mathcal{C}$ contains all nodes $j \in \mathcal{R}_s$ such that $\alpha_{jd} > 2$, while subset $\mathcal{D}$ contains those nodes $j \in \mathcal{R}_s$ with exponent α_{jd} smaller than or equal to 2. Please note that, due to social-aware forwarding rules, once a relay in $\mathcal{C}$ is picked, all subsequent relays will be also drawn from $\mathcal{C}$, since nodes in $\mathcal{C}$ are "closer" to the destination than those in $\mathcal{D}$. As far as convergence is concerned, in the worst case, all nodes in $\mathcal{D}$ are exploited before those in $\mathcal{C}$. So, if we set $n - 1$, i.e., the maximum number of intermediate hops allowed, to be greater than or equal to $|\mathcal{D}|$, we are sure that, even in the worst case, a relay in $\mathcal{C}$ is eventually selected. Since for relays in $\mathcal{C}$ convergence is guaranteed (in fact, $\alpha_{jd} > 2$, when $j \in \mathcal{C}$), the overall expected delay will converge.   □

6.2 Multi-Copy Schemes

Frequently, social-aware schemes are multi-copy. In the following we analyze whether using multiple copies can help the convergence of the expected delay when social-aware schemes are in use. The proofs of this section follow the same line of reasoning of the corresponding social-oblivious versions, once substituting $\mathcal{P}_i$ with $\mathcal{R}_i$. For this reason, in the following we omit them.

6.2.1 Two-Hop Forwarding

First, we focus on the m-copy 2-hop scheme. To this aim, we derive Theorem 8, which is in turn based on the following lemmas. Please note that in this case sequence π_i only contains nodes that belong to $\mathcal{R}_i$.

Lemma 4 (max_i^{sa}). When intermediate relays are selected according to sequence π_i, the source node is able to deliver at most max_i^{sa} copies with finite first hop expected delay, with max_i^{sa} being equal to the following:

$$max_i^{sa} = \arg\min_m \{f_{max}^{sa}(m, \pi_i) > 0\}, \qquad (4)$$

where $f_{max}^{sa}(m, \pi_i) = m + \sum_{z=m}^{|\mathcal{R}_s|} \alpha_z^{(i)} - (2 + |\mathcal{R}_s|)$ and $\alpha_z^{(i)}$ denotes the exponent α_{sj} of the zth node in sequence π_i.

Corollary 4. Quantity max_i^{sa} derived in Lemma 4 takes values in the interval $[max_{lo}^{sa}, max_{up}^{sa}]$. The upper and lower bound on max_i^{sa} (corresponding to the best and worst case for convergence) are reached when π_i corresponds to nodes in $\mathcal{R}_s$ encountered in increasing and decreasing order of α_{sj}, respectively.

Lemma 5 (min_i^{sa}). Assuming that intermediate relays are selected in the order specified by sequence π_i, the expected delay from intermediate relays to the destination d will converge if and only if there are at least min_i^{sa} intermediate relays, with min_i^{sa} being equal to the following:

$$min_i^{so} = \arg\min_m \{f_{min}^{sa}(m, \pi_i) > 0\}, \qquad (5)$$

where $f_{min}^{sa}(m, \pi_i) = \sum_{z=1}^m \alpha_z^{(i)} - (1 + m)$ and $\alpha_z^{(i)}$ denotes the exponent α_{jd} associated with the zth node in encounter sequence $\pi_i - \{d\}$.

Corollary 5. Quantity min_i^{sa} derived in Lemma 5 takes values in $[min_{lo}^{sa}, min_{up}^{sa}]$. The upper and lower bounds on min_i^{sa} (corresponding to the worst and best case for convergence) are reached when π_i corresponds to the sequence of nodes (belonging to $\mathcal{R}_s$) ordered in increasing and decreasing order of their exponents α_{jd}, respectively.

Lemmas 4 and 5 are the social-aware equivalent of Lemmas 2 and 3. Using their results, the following theorem about the 2-hop convergence can be derived, with $m < |\mathcal{R}_s|$.

Theorem 8 (m-copy 2-hop social-aware scheme). $E[D_{sd}]$ achieves convergence if and only if the following condition holds true:

C8 $m \geq min_{up}^{sa} \wedge max_i^{sa} \geq min_i^{sa}, \forall i \in \{1, \ldots, |\mathcal{R}_s^p|\}$,

where set $\mathcal{R}_s^p$ is the set of all permutations π_i for set $\mathcal{R}_s$.

Corollary 6. A sufficient condition for the convergence of the expected delay under the social-aware m-copy two-hop forwarding scheme in Theorem 4 is given by the following:

$$C8_{[s]}   m \geq min_{up}^{sa} \wedge min_{up}^{sa} \leq max_{lo}^{sa}.$$

Comparing the social-aware m-copy 2-hop with its social-oblivious counter part is not straightforward. In Section 7 we prove analytically that there is no clear winner between the two, and that either one or both can achieve convergence depending on the mobility scenario considered.

6.2.2 Multi-Hop Forwarding

Finally, in Theorem 9 we consider the most general case in which the source node generates m copies for the message

and each of them travel up to n hops along independent paths. We find that also in the social-aware case, multiple copies used together with multiple hops do not improve convergence with respect to the simple 1-copy n-hop scheme. Intuitively, this is because in the worst case (which occurs with non-negligible probability) all copies, after the first hop, follow the same multi-hop path, which requires conditions C6 and C7 for convergence.

Theorem 9 (m-copy n-hop social-aware scheme). $E[D_{sd}]$ *converges if and only if conditions C6 and C7 in Theorem 7 hold true.*

6.3 Discussion

Table 1 summarizes the convergence conditions for social-aware schemes derived so far. As in the social-oblivious case, multi-hop schemes do not benefit from the use of multiple copies, and in fact the 1-copy n-hop scheme and the m-copy n-hop scheme share the same convergence conditions. Similarly, the difference between 2-hop schemes mirrors that between the corresponding social-oblivious versions. Thus, the 1-copy 2-hop scheme is effective when $\alpha_{jd} > 2$ for all $j \in \mathcal{R}_s$, since it allows us to save resources by sending a single copy. However, when condition C5 does not hold, the only chance to achieve convergence is to exploit multiple copies.

If we focus on single-copy schemes, it is interesting to note that, differently from the social-oblivious case in which using additional hops did not provide any advantage, 1-copy social-aware schemes may benefit from multiple hops. In fact, for the 1-copy 2-hop scheme we need to impose that all intermediate relays j meet the destination with $\alpha_{jd} > 2$, which is a quite strong condition. On the other hand, if we use multiple hops (1-copy n-hop case), conditions C6 and C7 are required, which are milder than C5. More specifically, assuming that there are no limitations to the value that we can assign to n, condition C7 can be easily satisfied. Then, C6 relates to the convergence of the minimum of a set of Pareto random variables, which is always easier to achieve than for any single random variable from the set (corresponding to condition C5). The only constraint for the 1-copy n-hop case is that there must be at least one node z (the one with the highest fitness) meeting the destination with $\alpha_{zd} > 2$. In fact, for z, $\mathcal{R}_z = \{d\}$.

Finally, we compare the m-copy 2-hop case with the 1-copy n-hop case (which is equivalent to the m-copy n-hop scheme). There is no clear winner here, as each scheme can provide convergence when the other one cannot. For example, consider the case in which the source node is not able to send more than one copy (i.e, $max_i^{sa} = 1, \forall i \in \{1, \ldots, |\mathcal{R}_s^p|\}$). In this case, the m-copy 2-hop scheme becomes effectively a 1-copy 2-hop scheme, which fails to achieve convergence if some intermediate hop j does not have exponent α_{jd} greater than 2 (condition C5). Instead, exploiting multiple hops pays off in this case, as it allows us to rely on more intermediate relays, which may not meet the destination within a finite expected time but can bring the message "closer" to nodes that do meet d with $\alpha_{jd} > 2$. Vice versa, when $max_i^{sa} > 1$ for some i, the cooperative delivery of the multiple copies can overcome the presence of intermediate relays for which conditions C6-C7 do not hold. For example, when

there is not even one relay j with $\alpha_{jd} > 2$, then the m-copy 2-hop case is the only possible choice.

7 COMPARING SOCIAL-AWARE AND SOCIAL-OBLIVIOUS STRATEGIES

In the previous sections we have separately analyzed the convergence properties of social-oblivious and social-aware forwarding schemes, identifying the best strategies, from the convergence standpoint, for each of the two categories. In the following we take the champions of each class and we investigate whether there is a clear winner between social-oblivious and social-aware strategies when it comes to the convergence of their expected delay.

Let us first consider the case $\alpha_{sd} > 2$. We have seen in Section 5.3 that with this configuration the Direct Transmission scheme is the best choice from the convergence standpoint. In fact, with social-oblivious schemes using more than one hop, "bad" relays can be selected even starting from a source that is already able to reach the destination with a finite expected residual intermeeting time. This does not happen with social-aware strategies. In fact, assume that the source is the only node with $\alpha_{sd} = 2 + \varepsilon$, while all other nodes meet the destination with $\alpha_{jd} = 1 + \varepsilon$, with ε being a very small quantity. In the social-aware case, $\mathcal{R}_s$ contains only the destination, as all other nodes are clearly worse than the source node as relay. This shows the adaptability of social-aware schemes: the additional knowledge that they exploit makes them able to resort to simpler approaches (in this case, $\mathcal{R}_s = \{d\}$ is equivalent to the Direct Transmission) when they realize that additional resources in terms of number of copies or number of hops would not help the forwarding process. This implies that one can safely use the m-copy 2-hop or the 1-copy n-hop social-aware protocols because in the worst case they will do no harm (they will downgrade to simpler strategies, without exploiting wrong paths), while in the best case they are able to improve the convergence of the forwarding process.

When $\alpha_{sd} \leq 2$ and $\alpha_{jd} > 2$ for all nodes j in the relay set (i.e., $j \in \mathcal{R}_s - \{d\}$ for the social-aware case and $j \in \mathcal{P}_s - \{d\}$ for the social-oblivious case), the strategy of choice is the 1-copy 2-hop for both the social-oblivious and social-aware category. However, the 1-copy 2-hop social-aware scheme is overall more advantageous than its social-oblivious counterpart. More specifically, when the source node is the worst relay for the destination (i.e., $\min_i\{\alpha_{id}\} = \alpha_{sd}$), the social-oblivious and the social-aware approaches are equivalent (given that $\mathcal{P}_s = \mathcal{R}_s$). In all other cases, instead, $\mathcal{R}_s \subset \mathcal{P}_s$, thus, for the set of nodes in $\mathcal{P}_s - \mathcal{R}_s$, social-aware forwarding does not impose any constraint, while social-oblivious forwarding needs to impose constraints, thus resulting in stricter conditions for convergence.

Let us now focus on the remaining cases, namely i) when $\alpha_{sd} \leq 2$ and not all intermediate relays have exponent greater than 2, and ii) when $\alpha_{jd} \leq 2$ for all nodes j. In the first case, the social-aware m-copy 2-hop, the social-aware 1-copy n-hop, and the social-oblivious m-copy 2-hop can achieve convergence. In the second case, the only options for convergence are the social-aware m-copy 2-hop and the
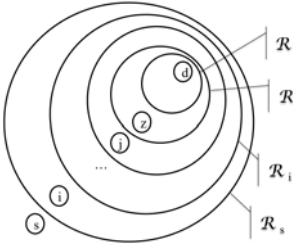
Fig. 1. Social forwarding at a glance.

$$\boldsymbol{\alpha} = \begin{pmatrix} 0 & 1.22 & 1.22 & 1.22 & 1.22 & 1.22 & 1.22 & 1.22 & 1.22 & 1.125 \\ 1.22 & 0 & 1.33 & 1.33 & 1.33 & 1.33 & 1.33 & 1.33 & 1.33 & 1.13 \\ 1.22 & 1.33 & 0 & 1.44 & 1.44 & 1.44 & 1.44 & 1.44 & 1.44 & 1.14 \\ 1.22 & 1.33 & 1.44 & 0 & 1.55 & 1.55 & 1.55 & 1.55 & 1.55 & 1.15 \\ 1.22 & 1.33 & 1.44 & 1.55 & 0 & 1.66 & 1.66 & 1.66 & 1.66 & 1.16 \\ 1.22 & 1.33 & 1.44 & 1.55 & 1.66 & 0 & 1.77 & 1.77 & 1.77 & 1.17 \\ 1.22 & 1.33 & 1.44 & 1.55 & 1.66 & 1.77 & 0 & 1.88 & 1.88 & 1.18 \\ 1.22 & 1.33 & 1.44 & 1.55 & 1.66 & 1.77 & 1.88 & 0 & 1.99 & 1.19 \\ 1.22 & 1.33 & 1.44 & 1.55 & 1.66 & 1.77 & 1.88 & 1.99 & 0 & 2.1 \\ 1.11 & 1.12 & 1.13 & 1.14 & 1.15 & 1.16 & 1.17 & 1.18 & 2.1 & 0 \end{pmatrix}$$

Fig. 2. Exponent matrix.

social-oblivious m-copy 2-hop. We first highlight the differences between the n-hop approach and the 2-hop approach by discussing when the social-aware 1-copy n-hop outperforms the other two strategies in terms of convergence (which can only happen in case i), then we focus on the social-aware and social-oblivious m-copy 2-hop strategies, thus covering both case i and ii.

So, assume that there exists at least one node z that meets the destination with $\alpha_{zd} > 2$. The m-copy 2-hop strategies send multiple copies to a set of relays, which in turn can only deliver the message to the destination directly. This implies that intermediate relays must have collectively the capability of reaching the destination, for all subsets with size m of possible relays. Here, only meetings with the destination are relevant, and if all relays but z have very low exponent for encounters with the destination, convergence may not be achieved. Differently from the 2-hop strategies, the social-aware n-hop scheme do not rely exclusively on the capabilities of meeting with d, but it is able to generate a *path* towards the destination in which intermediate nodes may not be good relays for d but good relays towards nodes with high fitness (in the extreme case, only $\alpha_{zd} > 2$ can hold). Thus, in the n-hop case, as long as the message can leave intermediate relays within a finite expected time, this could be enough for convergence. An example scenario is provided in Section 7.1. When all three strategies achieve convergence, the one to be preferred can be chosen based on resource consumption considerations. With the m-copy 2-hop strategies there can be up to $2m$ transmissions, while with the 1-copy n-hop scheme there are n. Hence, when $n < 2m$, the single-copy scheme should be preferred.

Let us finally compare the social-oblivious and the social-aware m-copy 2-hop schemes. Since they seem to cover similar mobility scenarios (as discussed in the previous section) and to be based on similar mechanisms (the min_i and max_i quantities, whose relation with m determines the convergence), it may be difficult to intuitively evaluate which one performs better in terms of convergence. For this reason, Theorem 10 below (whose proofs can be found in [25]) we tackle this problem from an analytical perspective. In [25, Appendix E, available in the online supplementary material] we provide a concrete example for both cases.

Theorem 10. *Since both the following configurations are feasible under the conditions in Lemma D1, it may happen that either*

the social-oblivious m-copy 2-hop scheme achieves convergence when the social-aware m-copy 2-hop scheme does not (Equation (6)), or vice versa (Equation (7)), depending on the underlying mobility process:

$$max_i^{so} \geq min_i^{so} \geq min_i^{sa} > max_i^{sa} \qquad (6)$$

$$min_i^{so} > max_i^{so} \geq max_i^{sa} \geq min_i^{sa}. \qquad (7)$$

Intuitively, an example of the first case is when there are a lot of nodes that meet the source with high α_{sj} (thus resulting in high max_i^{so}, high enough to be greater than min_i^{so}); if those relays have very low α_{jd}, they will not be used by the social-aware scheme, thus resulting in a low max_i^{sa}, possibly not high enough to guarantee that the second hop converges. It is easy to construct a corresponding example for the other case.

7.1 Example

In order to complement the theoretical discussion of the previous section, in the following we provide a concrete example for the case in which the social-aware 1-copy n-hop scheme is the only one achieving convergence. In [25] we also provide two concrete examples for the cases discussed in Theorem 10.

The mobility scenario we consider is described by the exponent matrix in Fig. 2. Please note that such matrix has been chosen just to exemplify the behaviour of the social-aware 1-copy n-hop scheme. For a more realistic analysis, please refer to Section 8. Element α_{ij} in matrix $\boldsymbol{\alpha}$ (of size 10) gives the Pareto exponent for the i, j node pair. We assume that node $i = 1$ is the source node and that node $j = 10$ is the destination. In this case the source node is the node with the lowest fitness value, thus the m-copy 2-hop social-oblivious and social-aware schemes overlap (in fact, $\mathcal{P}_s = \mathcal{R}_s$).

We start with the 1-copy n-hop scheme. The size of set $\mathcal{D}$ is 8, since there are eight nodes with $\alpha_{jd} \leq 2$. Thus, we need to set the maximum number of allowed hop n to 9 (condition C7). Then, we compute $\sum_{j \in \mathcal{R}_i} \alpha_{ij} - (1 + \mathcal{R}_i)$ (condition C6) for all nodes $i \in \mathcal{R}_s \cup \{s\}$ (Table 2). Since the computed quantities are greater than zero for all possible relays (including the source node), the 1-copy n-hop social-aware scheme achieves convergence in this scenario.

We now focus on the m-copy 2-hop scheme, recalling that the social-oblivious version and the social-aware

TABLE 2
Condition C6 for Each Relay (Including the Source)

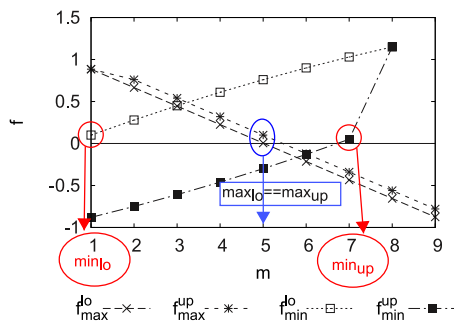|	1	2	3	4	5	6	7	8	9
$C6$	0.89	1.44	1.78	1.9	1.8	1.48	0.94	0.18	0.1

Fig. 3. min_i and max_i for the m-copy 2-hop scheme.

version are equivalent in this case (thus we drop superscripts *so* and *sa*). In order to verify sufficient condition C3$_{[s]}$, we need to find max_{lo} and min_{up}, i.e., the value of max_i and min_i in the worst case. According to Corollary 1, max_{lo} is reached when permutation π_i corresponds to relays encountered in decreasing order of α_{sj}, while, according to Corollary 2, min_{up} is achieved when permutation π_i corresponds to relays encountered by the source node in increasing order of α_{jd}. We denote these two permutations as π_i^* and π_i' respectively. In Fig. 3, we plot function $f_{max}^{lo}(m) = f_{max}(m, \pi_i^*)$ corresponding to the case in which max_{lo} is reached and function $f_{min}^{up}(m) = f_{min}(m, \pi_i')$ corresponding to the case in which min_{up} is achieved. Recall that min_{up} corresponds to the first m value for which f_{min}^{up} is greater than zero, thus $min_{up} = 7$. Similarly, max_{lo} corresponds to the last m value for which function f_{max}^{lo} is greater than zero, and so $max_{lo} = 5$. Since $max_{lo} < min_{up}$, sufficient condition C3$_{[s]}$ is not satisfied.

It is easy to show that also the necessary and sufficient condition C3 does not hold. Recall that the necessary and sufficient condition states that convergence is ensured as long as $max_i \geq min_i$ for all encounter permutations π_i. However, this does not happen here. Consider (Fig. 3) f_{max}^{up} and f_{min}^{lo}, i.e., functions f_{max} and f_{min} in the best case. The first integer values of m before the functions become negative determine the values of max_{up} and min_{lo}. Since $max_{up} = 5$, from Corollary 1 we have that max_i varies in the range $[5, 5]$, i.e., max_i is always equal to 5 regardless of the permutation considered. This means that, for the permutation corresponding to the min_i worst case, the source node will not be able in any case to send more than five copies with finite first-hop expected delay (while seven are required). Hence, convergence cannot be achieved in this case.

8 CONVERGENCE IN MOBILITY TRACES

We conclude the paper by applying the convergence conditions derived in the previous sections to a set of contact traces that are frequently used in the literature: Cambridge [26], Infocom'05 [26], and RollerNet [27], composed respectively of 12, 41, and 62 nodes. Due to space limitations, we do not recall further properties of these traces and we move straight to the application of our convergence conditions. The Pareto exponents are estimated using Maximum Likelihood Estimation, setting t_{min} equal to the sampling interval used for collecting the trace ($120s$ for both Infocom'05 and Cambridge, $15s$ for RollerNet). Under this configuration,

after applying the Cramer-von Mises criterion, we obtain that the Pareto hypothesis cannot be rejected for 80, 97, and 85.5 percent of pairs for Cambridge, Infocom'05, and RollerNet, respectively.[7]

In Fig. 4 we show the percentage of pairs for which the strategies identified in Section 4 achieve convergence, in both the social-oblivious and the social-aware case. We omit the n-hop schemes that share the convergence conditions of other strategies that consumes less resources, as discussed in Sections 5.3 and 6.3. However, we provide the results obtained by applying the sufficient convergence conditions in Corollaries 3 and 6, in order to highlight that in these cases they identify correctly the right set of pairs. In the Cambridge data set (Fig. 4a), the best performance in terms of convergence is delivered by the social-oblivious m-copy 2-hop scheme, with the social-aware m-copy 2-hop scheme just slightly behind. In the Infocom data set (Fig. 4b) there is basically a tie between the same two strategies, with the social-aware one performing slightly better in this case. In the RollerNet scenario (Fig. 4c), the social-aware m-copy 2-hop scheme clearly outperforms the others.

From the Pareto exponent distribution point of view, the Cambridge data set is the one more shifted towards higher values ($\min = 1.39, \text{median} = 1.68, \max = 2.86$) with respect to Infocom ($\min = 1.24, \text{median} = 1.43, \max = 1.82$) and RollerNet ($\min = 1.27, \text{median} = 1.58, \max = 3.27$). This is reflected by the performance of the Direct Transmission scheme. In the Infocom data set no pair meets with an exponent higher than 2. This is a very unfortunate case from the convergence standpoint, but the m-copy 2-hop schemes are still able to overcome this limitation and to reach around 80 percent of pairs with a convergent expected delay. Fig. 4 also confirms the importance of taking into account heterogeneity when studying the convergence of forwarding strategies. In fact, applying the convergence condition for the homogeneous case derived in [25] ($\alpha > \frac{2}{N} + 1$, assuming that we set $m > \frac{N}{2}$) and using the estimated exponent for the aggregate intermeeting times ($\alpha = 1.44$), we would have obtained a "convergent" verdict for the m-copy 2-hop scheme *for all source-destination pairs* in the Infocom scenario, but this is not always the case, as shown in Fig. 4b.

9 CONCLUSIONS

Assuming heterogenous Pareto intermeeting times, in this paper we have derived the conditions on the Pareto exponents such that the expected delay of a large family of forwarding protocols is finite. Our main result for the social-oblivious case is that convergence is not improved by using more than two hops (and in some conditions direct transmission, with just one hop, is the most efficient choice). In the social-aware case, instead, allowing more than two hops can provide convergence when all other strategies fail. As for the comparison of single-copy and multi-copy schemes, we found that multi-copy strategies can, in some cases, outperform single-copy strategies in terms of convergence of the expected delay. The use of multiple copies, in fact,

---

7. Different acceptance percentages in [27] are due to a different t_{min} setting. We believe that it is more correct to set t_{min} equal to the granularity of the trace because samples smaller than the granularity are a just a few and only emerge due to statistical fluctuations.
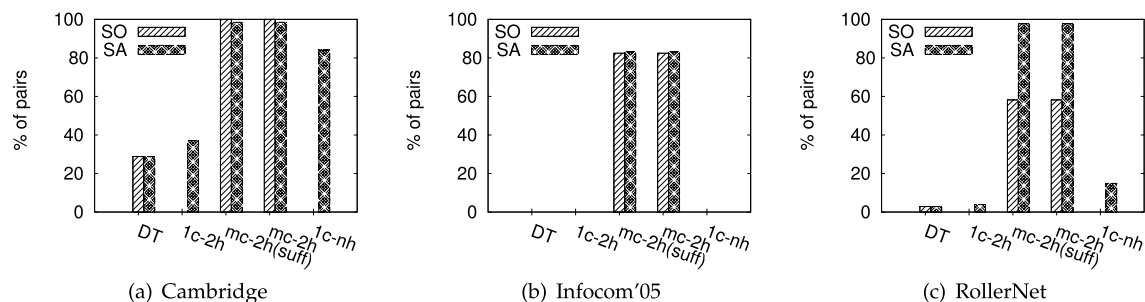
Fig. 4. Percentage of convergent pairs in traces.

benefits from the parallel delivery of the message from different nodes, which may overcome the limitations of individual nodes in achieving a finite expected delay. Finally, comparing social-oblivious and social-aware multi-copy solutions we were able to prove mathematically that there is not a clear winner between the two, since either one can achieve convergence when the other fails depending on the underlying mobility scenario.

Besides the theoretical value of the above convergence model per se, we believe that such model has also important practical implications. For the majority of forwarding schemes, nodes would be able to evaluate online whether a policy can achieve convergence or not (hence they can decide which one is to be preferred). For example, convergence can be easily verified for 1-hop and 2-hop strategies, since it is perfectly reasonable to assume that nodes can learn the Pareto exponents of their direct neighbours. The only relevant policy (from the convergence standpoint) for which nodes may not be able to verify the convergence online (because it would require the knowledge of the exponents of potentially distant nodes) is the social-aware m-copy n-hop scheme. One way to address this problem is to let the source node pick this strategy only as a last resort, i.e., only when it is not possible to collect the required exponents for testing its convergence and when two-hop schemes are not able to achieve it.

ACKNOWLEDGMENTS

REFERENCES

[1]   T. Spyropoulos, K. Psounis, and C. Raghavendra, "Efficient routing in intermittently connected mobile networks: The single copy case," *IEEE/ACM Trans. Netw.*, vol. 16, no. 1, pp. 63–76, Feb. 2008.

[2]   M. Grossglauser and D. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Trans. Netw.*, vol. 10, no. 4, pp. 477–486, Aug. 2002.

[3]   T. Spyropoulos, K. Psounis, and C. Raghavendra, "Efficient routing in intermittently connected mobile networks: The multiple-copy case," *IEEE/ACM Trans. Netw.*, vol. 16, no. 1, pp. 77–90, Feb. 2008.

[4]   C. Lee, "Heterogeneity in contact dynamics: Helpful or harmful to forwarding algorithms in DTNs?" in *Proc. Int. Conf. Model. Optimization Mobile, Ad Hoc, Wireless Netw.*, 2009, pp. 1–10.

[5]   T. Spyropoulos, T. Turletti, and K. Obraczka, "Routing in delay-tolerant networks comprising heterogeneous node populations," *IEEE Trans. Mobile Comput.*, vol. 8, no. 8, pp. 1132–1147, Aug. 2009.

[6]   A. Picu, T. Spyropoulos, and T. Hossmann, "An analysis of the information spreading delay in heterogeneous mobility DTNs," in *Proc. IEEE Int. Symp. World Wireless, Mobile Multimedia Netw.*, 2012, pp. 1–10.

[7]   Z. Haas and T. Small, "A new networking model for biological applications of ad hoc sensor networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 1, pp. 27–40, Feb. 2006.

[8]   A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on opportunistic forwarding algorithms," *IEEE Trans. Mobile Comput.*, vol. 6, no. 6, pp. 606–620, Jun. 2007.

[9]   H. Cai and D. Eun, "Crossing over the bounded domain: From exponential to power-law intermeeting time in mobile ad hoc networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 5, pp. 1578–1591, Oct. 2009.

[10]   T. Karagiannis, J.-Y. Le Boudec, and M. Vojnovic, "Power law and exponential decay of intercontact times between mobile devices," *IEEE Trans. Mobile Comput.*, vol. 9, no. 10, pp. 1377–1390, Oct. 2010.

[11]   A. Passarella and M. Conti, "Analysis of individual pair and aggregate inter-contact times in heterogeneous opportunistic networks," *IEEE Trans. Mobile Comput.*, vol. 12, no. 12, pp. 2483–2495, Dec. 2013.

[12]   V. Conan, J. Leguay, and T. Friedman, "Characterizing pairwise inter-contact patterns in delay tolerant networks," in *Proc. 1st Int. Conf. Autonomic Comput. Commun. Syst.*, pp. 1–9, 2007.

[13]   C. Boldrini, M. Conti, and A. Passarella, "Performance modelling of opportunistic forwarding under heterogenous mobility," *Comput. Commun.*, vol. 48, pp. 56–70, 2014.

[14]   C. Lee and D. Eun, "Exploiting heterogeneity in mobile opportunistic networks: An Analytic Approach," in *Proc. IEEE 7th Annu. Commun. Soc. Conf. Sensor Mesh Ad Hoc Commun. Netw.*, 2010, pp. 1–9.

[15]   A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Durham, NC, USA, Tech. Rep. CS-200006, Duke University.

[16]   A. Lindgren, A. Doria, and O. Schelén, "Probabilistic routing in intermittently connected networks," in *Proc. Workshop Serv. Assurance Partial Intermittent Resources*, 2004, pp. 239–254.

[17]   P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: Social-based forwarding in delay tolerant networks," *IEEE Trans. Mobile Comput.*, vol. 10, no. 11, pp. 1576–1589, Nov. 2011.

[18]   E. Daly and M. Haahr, "Social network analysis for information flow in disconnected Delay-Tolerant MANETs," *IEEE Trans. Mobile Comput.*, vol. 8, no. 5, pp. 606–621, May. 2009.

[19]   C. Boldrini, M. Conti, and A. Passarella, "Exploiting users' social relations to forward data in opportunistic networks: The HiBOp solution," *Perv. Mob. Comput.*, vol. 4, no. 5, pp. 633–657, 2008.

[20]   P. Costa, C. Mascolo, M. Musolesi, and G. Picco, "Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 5, pp. 748–760, Jun. 2008.

[21]   C. Boldrini, M. Conti, and A. Passarella, "Less is more: Long paths do not help the convergence of social-oblivious forwarding in opportunistic networks," *ACM/SIGMOBILE MobiOpp*, vol. 12, pp. 1–8, 2012.

[22]   S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: an approach to interplanetary internet," *IEEE Commun. Mag.*, vol. 41, no. 6, pp. 128–136, Jun. 2003.

[23]   C. Boldrini, M. Conti, and A. Passarella, "From pareto inter-contact times to residuals," *IEEE Commun. Lett.*, vol. 15, no. 11, pp. 1256–1258, Nov. 2011.

[24] N. Johnson and S. Kotz, *Distributions in Statistics: Continuous Univariate Distributions: Vol. 1.* New York, NY, USA: Houghton Mifflin, 1970.

[25] C. Boldrini, M. Conti, and A. Passarella. (2013). The stability region of the delay in Pareto opportunistic networks," IIT-CNR, Tech. Rep. TR-13/2013, [Online]. Available: http://cnd.iit.cnr.it/chiara/pub/techrep/boldrini2013convergence_tr.pdf

[26] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proc. ACM SIGCOMM Workshop Delay-Tolerant Netw.*, 2005, pp. 244–251.

[27] P.-U. Tournoux, J. Leguay, F. Benbadis, J. Whitbeck, V. Conan, and M. D. de Amorim, "Density-aware routing in highly dynamic DTNs: The rollernet case," *IEEE Trans. Mobile Comput.*, vol. 10, no. 12, pp. 1755–1768, Dec. 2011.

Chiara Boldrini received the MSc degree in computer engineering and the PhD degree in information engineering both from the University of Pisa, Italy, in 2006 and 2010, respectively. She is a researcher at IIT-CNR. Her research interests are in the area of opportunistic and delay-tolerant networking. More specifically, she has been studying routing, content dissemination, and energy saving issues, with a special focus on performance modelling using stochastic modeling techniques and complex network analysis. Additionally, she has been working on models for human mobility that can realistically represent the behaviour of mobile users.

Marco Conti is a research director of the Italian National Research Council (CNR). He has published in journals and conference proceedings more than 300 research papers and four books related to design, modeling, and performance evaluation of computer networks, pervasive systems and social networks. He is an editor-in-chief of *Elsevier Computer Communications journal* and an associate editor-in-chief of *Elsevier Pervasive and Mobile Computing*. He received the Best Paper Award at IEEE ISCC 2012 and IEEE WoWMoM 2013. He has served as a General and Program (co)chair for several conferences, including Networking 2002, IEEE WoWMoM 2005 and 2006, IEEE PerCom 2006 and 2010, ACM MobiHoc 2006, and IEEE MASS 2007. He is a member of the IEEE.

Andrea Passarella received the PhD degree in computer engineering in 2005. He is with IIT-CNR, Italy. He was a research associate at the Computer Laboratory, Cambridge, United Kingdom. He published 100+ papers on mobile social networks, opportunistic, ad hoc and sensor networks, receiving the best paper award at IFIP Networking 2011 and IEEE WoWMoM 2013. He was the PC co-chair of IEEE WoWMoM 2011, Workshops co-chair of IEEE PerCom and WoWMom 2010, and the co-chair of several IEEE and ACM workshops. He is in the Editorial Board of *Elsevier Pervasive and Mobile Computing* and Inderscience *IJAACS*. He was the guest co-editor of several special sections in ACM and Elsevier Journals. He is the vice-chair of the IFIP WG 6.3 "Performance of Communication Systems." He is a member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.