Consiglio Nazionale delle Ricerche

The stability region of the delay in Pareto opportunistic networks

C. Boldrini, M. Conti, A. Passarella

IIT TR-13/2013

Technical report

Luglio 2013

Istituto di Informatica e Telematica

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. Finally, we 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.

Index Terms—Wireless communication, Routing protocols, Mobile Computing, Algorithm/protocol design and analysis, Mobile communication systems, Ubiquitous computing, Modelling techniques

✦

1 INTRODUCTION

THE 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. On the one hand, there are *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 (avoiding at most those nodes that have already forwarded the message). The main advantage of these strategies is that they are intrinsically simple and lightweight (practically no

information to collect, store, or mine). This simplicity, however, is typically paid in terms of suboptimal routing performance. In order to improve message forwarding, smarter strategies have been proposed that exploit information on the social context users operate in. These approaches, referred to as *social-aware*[1], 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. In the first case, at any time, in the network there is just one copy of the message to be delivered, while in the second case more copies are generated, hoping that at least one of them will eventually reach the destination. Multi-copy strategies have been shown to improve the reliability of delivery with respect to single-copy approaches [2]. 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 [3]), 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. Since messages follow multi-hop paths across the nodes of the network, their delay is the result of the delay accumulated at

- *C.Boldrini, M.Conti, and A.Passarella 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

1. These policies are also referred to as utility-based [1], in contrast to randomized strategies, corresponding to our social-oblivious schemes. In the following we will stick to the social-oblivious vs social-aware classification, in order to highlight the fact that the forwarding utility is almost always computed from mobility and social information.

each hop along the forwarding path. Therefore, the time (*intermeeting time*) between consecutive encounters of a pair of nodes is the elementary component of the overall delay. Thus, 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] [2] [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 also with non exponential intermeeting times, as shown in this paper. In some cases, the distribution of intermeeting times can drastically affect the convergence of the expected delay. This happens, for example, when intermeeting times feature a Pareto[2] 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. Being the delay the result of the composition of the time intervals between node encounters, depending on the exponent values featured by intermeeting times, the expectation of the delay itself might diverge.

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]. Until then, intermeeting times were assumed to feature i.i.d. exponential distributions. The hypothesis in [8] was later revised by [9] in favor of Pareto with exponential cut-off intermeeting times, but this finding has been also questioned [10] [11]. Despite these contradicting results, we believe that Pareto intermeeting times are a case worth considering for three main reasons. First, a recent analysis of pairwise intermeeting times [11] (considering both one of the traces used by [8] and a new trace) has confirmed that the Pareto hypothesis for intermeeting times cannot be ruled out in general. Second, Cai and Eun [12] 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. Third, the Pareto distribution is one of the most popular representa-

tives of high-variance distributions. These distributions form a class opposite to the exponential distribution in that, due to their high variability, they can have a divergent expectation, which in turn can bring to delay convergence issues. Since, based on the trace analyses available in the literature there is mounting request for models where intermeeting times are not exponential [4], and because it is not possible to rule out high-variance distributions from the set of plausible distributions for intermeeting times [13], we believe it is of paramount importance to evaluate how the achieved performance of networking protocols change when the assumption of exponential (or, in general, low-variance) intermeeting times is released and, due to their representativeness, in the case of Pareto intermeeting times.

Under the Pareto intermeeting times assumption, in this paper 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. Homogeneous mobility implies that the intermeeting times between any pair of nodes have the same statistical characteristics (e.g., same exponent for Pareto intermeeting times). However, measurement studies [14] [8] have shown that real networks are intrinsically heterogeneous. Thus, in this paper, we assume heterogenous pairwise intermeeting times and we investigate whether heterogeneity in contact patterns helps the convergence of the expected delay of a general class of social-oblivious and social-aware 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. 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. Intuitively, since social-oblivious protocols are by definition not able to select increasingly better forwarders (while social-aware strategies are, instead), nodes picked at the first hop are statistically the same as those picked at the n-th hop, thus two hops are enough for exploring the diversity available in the network and to avoid that messages get stuck at the source node.
- 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. The reason is that only n-hop social-aware schemes are able to establish a multi-hop path from the source node to the destination node, along which the chances of encountering the destination are always increasing.

- 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. This is due to the fact that a parallel delivery of more than one copy can increase the chances of reaching the destination.

- Finally, 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.

The paper is organised as follows. In Section 2 we briefly review the state of the art on forwarding protocols for opportunistic networks. In Section 3 we describe the network model we consider and the assumptions we make. Then, in Section 4 we identify a set of representative classes of social-oblivious and social-aware schemes and, for these classes, we derive in Sections 5 and 6 the conditions for the expectation of their delay to be finite. In Section 7 these conditions are compared with each other, detecting the cases in which each of them performs better. Finally, Section 8 concludes the paper.

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. Most existing models assume that intermeeting times are approximately exponentially distributed [5] [6] [15], 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 [3], 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 [16] 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 an ideal scenario without resource limitations Epidemic achieves the minimum possible delay, but in realistic settings it is typically impractical due to the huge amount of resources it consumes [2]. 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 [2], gossiping [7]). Another popular social-oblivious forwarding protocol is the Two Hop scheme [3], 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 [3].

Social-aware strategies can have different levels of awareness. Simplest approaches exploit information such as time since the last encounter (Spray&Focus [2]) or frequency of encounters (PROPHET [17]). 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 [18], SimBet [19]). Alternatively, as in the case of HiBOp [20] and SocialCast [21], 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 [22], 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 [22], we include the analysis of social-aware forwarding strategies (Section 6) and a detailed comparison between social-aware and social-oblivious strategies from the convergence standpoint (Section 7).

3 NETWORK MODEL AND GENERAL RESULTS ABOUT PARETO INTERMEETING TIMES

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 [23], 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 assume that the transmission of a message can always be completed during a pairwise contact, 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 follows.

Definition 1 (Intermeeting Time): The intermeeting time M_{ij} between node i and node j is defined as the time between two consecutive meetings between the same pair of nodes[3].

For the sake of tractability, we assume that intermeeting times between every specific node pair i, j are independent and identically distributed and that their expectation $E[M_{ij}]$ does not vary with time (in other words, we assume a stationary network). By definition, the rate of encounter between node i and node j is given by $\frac{1}{E[M_{ij}]}$. Since intermeeting times between the same pair of nodes i, j are i.i.d., the encounter process between two nodes i and j can be seen as a renewal process with renewal intervals distributed as M_{ij} [24].

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 (Definition 2).

Definition 2 (Residual Intermeeting Time): Assuming that node i and node j are not in contact at a random time t_0, the residual intermeeting time R_{ij} between them is given by the time interval between t_0 and the first time node i and node j come into each other's range again.

The interplay between intermeeting times and residual intermeeting times in a forwarding model is not trivial. While it is straightforward to take the residual intermeeting time as the time between message generation (hereafter denoted as t_0) and relaying at the source node, it is less clear how to model the time interval before message handover when considering an intermediate relay i. In Lemma 1 at the end of the section we discuss how to model this time interval and we derive a worst-case result holding for Pareto intermeeting times.

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 following CCDF:

---

3. Without loss of generality, here we assume a deterministic unit disk graph model for radio propagation. In other words, nodes can communicate only if their current distance is smaller than the transmission range. This is a common assumption in the literature on opportunistic networks. The proposed framework still applies for every other model of radio propagation.

$$F_{M_{ij}}(t) = \left(\frac{t_{min_{ij}}}{t + t_{min_{ij}}} \right)^{\alpha_{ij}}, \qquad (1)$$

in which we use the definition of the Pareto distribution which allows for values arbitrarily close to zero, usually denoted as American Pareto [25] [26] (as opposed to the European Pareto version, where $F_{M_{ij}}(t) = (\frac{t_{min_{ij}}}{t})^{\alpha_{ij}}$) or Pareto distribution of the second kind [27]. Parameters α_{ij} and $t_{min_{ij}}$ are usually referred to as the shape and scale of the Pareto distribution, respectively. Note that we do not require intermeeting times M_{ij} and M_{ji} to be symmetric. Please note also that being the American Pareto a European Pareto shifted by $t_{min_{ij}}$ to the left, both Pareto definitions share the same requirements for their expectation to converge (Remark A1 in Appendix A). Similarly to the reference literature [8][9], in the following we restrict to the case of power law random variables having the same scale, i.e., $t_{min_{ij}} = t_{min}, \forall i, j$. The following remarks summarize properties of the Pareto distribution that will be used throughout Sections 5 and 6.

Remark 1: The Pareto distributions introduced above are defined for $\alpha_{ij} > 0$ (due to the required PDF normalization [29]), and their expectation converges (i.e., is finite) when $\alpha_{ij} > 1$.

As we have already discussed, residual intermeeting times come into the picture more often than intermeeting times, because the time of the generation of new messages can be modelled as a random time with respect to the evolution of the mobility process. Following a standard approach, in [26] we have shown that, from an American Pareto random variable with shape α_{ij} and scale $t_{min_{ij}}$, we obtain residuals that feature an American Pareto distribution with shape $\alpha_{ij} - 1$ and scale $t_{min_{ij}}$. In the case of European Pareto, the residual is not exactly Pareto distributed but it converges to a Pareto distribution with shape $\alpha_{ij} - 1$ in the tail [26]. Thus, it shares the same convergence conditions as the residual of an American Pareto random variable. For the residual intermeeting time, the following remark holds.

Remark 2: The Pareto distribution of R_{ij} is defined for $\alpha_{ij} > 1$ (due to the required PDF normalization), and its expectation converges when $\alpha_{ij} > 2$.

It is possible to prove (see Remark A2 in Appendix A) that conditioning does not affect convergence for Pareto random variables. More specifically, a Pareto random variable X conditioned to be greater than a constant value t_c (denoted as X^{t_c}) features a finite expectation under the same conditions described in Remark 1. This is an important property of the Pareto distribution that allows us to simplify, without loss of accuracy, the analysis of the convergence of the expected delay. In fact, in all cases in which we should consider intermeeting or residual intermeeting times conditioned to be greater than a certain value, we can simply take into account the unconditioned intermeeting or residual intermeeting time. These cases are discussed in detail in the Supplemental Material, while will not be further discussed here due to lack of space.

As anticipated a few paragraphs above, it is not straightforward to model the time interval before message hand-over to a generic node j when considering an intermediate relay i receiving its copy of the message at time t_i. In particular, assuming that j is a possible next hop under the forwarding strategy in use, the time before node i hands over the message to node j depends on whether nodes i and j met in time interval $[t_0, t_i]$. In fact, meetings correspond to renewals in the encounter renewal process between i and j, hence, from the meeting time on, we should consider the intermeeting time and not its residual (for more details, see the proof of Lemma 1 in Appendix B). 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). The reason is that this case corresponds to the worst case from the convergence standpoint. Since convergence must be enforced in all cases, if we focus on the worst case it is guaranteed that convergence in all other cases will follow. We will use the result in Lemma 1 throughout the mathematical analysis in Sections 5 and 6.

Lemma 1: 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 $R_{ij}^{t_i - t_0}$ (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 three main strategies that forwarding protocols can adopt in order to improve their forwarding performance, namely the number of hops allowed, the number of copies generated, and whether the source and relay nodes keep track of the evolution of the forwarding process or not. 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.

First, forwarding strategies can be single-copy or multi-copy. In the former case, at any point in time there can be at most one copy of each message circulating in the network. In the latter, multiple copies can travel in parallel, thus in principle multiplying the opportunities to reach the destination. 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 [2]) are left as future work.

Second, forwarding protocols can be classified based on the number of hops that they allow messages to traverse. In principle, this number could also be infinite. However, being such an approach not feasible in practice, the number of hops is either limited arbitrarily

(e.g., using the TTL field) or is 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.

Third, the amount of knowledge that each node in the forwarding process can rely on (or is willing to collect and store) is an additional element for classifying forwarding strategies. In this paper we only consider the case in which both the source node and the 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 nodes 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 will make protocols not scalable).

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 [3], in which the source node can only deliver the messages to the destination. The 2-hop 1-copy forwarding is equivalent to the Two Hop forwarding introduced in [3]. 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 number of hops are also defined as parameters of the algorithm.

4.1 Abstracting social-aware strategies

Due to the variety of social-aware schemes (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 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

---

4. In [22] 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. Note that we do not consider the memoryless case for social-aware strategies, as it is not logical to assume lack of memory for protocols that already keep track of other types of information on the network.

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 [30].

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. 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. Since we assume, for the sake of comparison with [8], that the probability of an encounter between any pair of nodes is strictly greater than zero, we have that $|\mathcal{P}_i| = N - 1$ for all nodes i. Since it is easy to show that, when $\alpha_{ij} \leq 1$, none of the forwarding algorithms considered in this paper are able to achieve a convergent expected delay (Corollary D1 in Appendix D), in the rest of the paper we only consider the case $\alpha_{ij} > 1$ for all i, j node pairs (which implies that the residual intermeeting times are defined, as discussed in Remark 2).

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 Remark 2 in Section 3, and we move to the analysis of the 1-copy 2-hop scheme immediately.

Theorem 1 (1-copy 1-hop scheme): When the 1-copy 1-hop relaying protocol is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if $\alpha_{sd} > 2$.

Theorem 2 (1-copy 2-hop scheme): When the 1-copy 2-hop relaying protocol is used, the expected delay for messages generated by the source node s for the destination node d 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 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 if 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 Lemma A4 in Appendix A, we know that $\min_{j \in \mathcal{P}_s}\{R_{sj}\}$ features a Pareto distribution with shape $\sum_{j \in \mathcal{P}_s}(\alpha_{sj} - 1)$, which, according to Remark 1, should be greater than 1 in order to have finite expectation. This implies $\sum_{j \in \mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$, thus obtaining condition C1. We now consider the convergence of the second hop. 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 if 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 if $\alpha_{jd} > 2$. Given that node j can be any node apart from s and d, condition $\alpha_{jd} > 2$ must hold for all nodes j different from s and d, and thus sufficient condition C2 is proved.

Conditions C1 and C2 are not only sufficient but also necessary conditions for the expected delay to be finite. In fact, if condition C1 is not satisfied a message can never leave its source node within a finite expected time, and thus its overall expected delay will not converge. Analogously, if condition C2 is not satisfied, there exists at least one relay that delivers the message to the destination with an infinite expected time, thus there is at least one two-hop divergent path. Since the 1-copy 2-hop scheme cannot control whether to choose a one-hop or a two-hop path nor avoid divergent relays, convergence at both the first hop and second hop (through all possible relays) is necessary for convergence. 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 the Supplemental Material. □

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): When the single-copy n-hop relaying protocol is used, the expected delay for messages generated by the source node s

for the destination node d converges if and only if conditions C1 and C2 in Theorem 2 hold true.

Proof: See Appendix C 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 enough for convergence. $\square$

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 provide at most an intuitive sketch for the proofs, which can be found in a detailed version in Appendix C.

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 the following we derive the convergence conditions for this case. 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. In Lemma 2, assuming that m is unbounded, we study what is the maximum number of copies with convergent first hop expected delay that the source node is able to hand over under social-oblivious forwarding. Please note that it is possible to prove that first-hop convergence becomes more difficult as the number of available relays decreases. The number of available relays decreases as the source node hands over copies, since relays cannot be used twice. Hence, after a certain point, the number of relays left does not allow convergence to be achieved, 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 i-th 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 the multi-copy social-oblivious two-hop forwarding protocol is in use and 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:

$$max_i^{so} = \arg\min_m \{f_{max}^{so}(m, \pi_i) > 0\}, \tag{2}$$

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 z-th node belonging to π_i.

Proof: Since the source is memoryful, after the k-th copy is relayed, the next copy can be delivered only to the subset of nodes that comprises only those that have not been already used as relay. It is possible to prove (Lemma C1 in Appendix C) that the convergence conditions become stricter as the cardinality of the set from which we choose the relays decreases. This lets us focus on the delivery of the m-th copy, because that is the one that sees the smallest set of possible relays, whose cardinality is $N-m$. Similarly to the line of reasoning used in the proof of Theorem 2, the m-th copy is relayed with finite delay if the minimum of the residuals between the source and the possible relays is finite, which results in a condition on the sum of the exponents of the corresponding intermeeting time distributions. Specifically, once we focus on a specific sequence π_i of encounters at the source node (we prove that any of these sequences can happen under our assumptions), we derive that the m-th copy is relayed within a finite expected time if the sum of the $N-m$ exponents α_{sj} associated with the last $N-m$ nodes in π_i is greater than $|\mathcal{P}_s|+2-m$. The latter result can be used to derive (Equation 2) the maximum m value under which this condition remains true, corresponding to max_i^{so}. In fact, since the convergence condition worsens when m increases, there is a cut-off m beyond which first hop convergence is never achieved for a given sequence of encounters π_i and it corresponds to the highest integer m value for which the above condition is still true. $\square$

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 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 (Lemma C1 in Appendix C) 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}): Under multi-copy social-oblivious two-hop forwarding, 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:

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

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 z-th 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 i-th relay, assuming that it receives its copy of the message at time t_i. The time before the i-th 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 z-th node in π_i. Then, applying Remark 1 to the minimum of the residuals, 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 Appendix C. □

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.

Theorem 4 (m-copy 2-hop scheme): When the m-copy two-hop forwarding protocol is used (with $m < N - 1$), the expected delay for messages generated by the source node s for the destination node d converges if and only if the following condition holds true:

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

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 implies 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[5] 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-

---

5. 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.

hop forwarding scheme in Theorem 4 is given by the following:

C3$_{[s]}$ $m \geq min_{up}^{so} \wedge min_{up}^{so} \leq max_{lo}^{so}$.

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 Appendix C.2.1 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 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): When the social-oblivious m-copy n-hop protocol is used, the expected delay for messages generated by the source node s for the destination node d 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 Appendix C 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. 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, both n-hop social-oblivious protocols and the 1-copy 2-hop scheme share the same convergence conditions (C1 and C2), but the former consumes much more network resources than the latter. This tells us that, 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. This is confirmed by the fact that 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 is not able to directly deliver the message within a finite expected time, and thus exploring more relays is convenient as it allows the source node to exploit node diversity. In fact, even if the source node cannot reach destination d directly with a finite expected delay, it may be able to hand over the message to other nodes within a finite expected time. 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.

When there exists at least one intermediate relay which is not able to deliver the message directly to the destination within a finite expected time, the most effective strategy is the m-copy 2-hop forwarding. In fact, with m-copy 2-hop forwarding 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. 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

	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]

TABLE 1
Summary of convergence conditions for social-oblivious and social-aware routing strategies

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): When the single-copy social-aware 2-hop forwarding strategy is used, the expected delay for messages generated by the source node s for the destination node d 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. $\square$

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): When the social-aware 1-copy n-hop forwarding strategy is used, the expected delay for messages generated by the source node s for the destination node d 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 $\{s, i, \cdots, j, z, d\}$ with increasing fitness (Figure 1), the relation $\mathcal{R}_s \supset \mathcal{R}_i \supset \cdots \mathcal{R}_j \supset \mathcal{R}_z$ holds. Exploiting Lemma 1, we know that the time before the message leaves a generic node i is distributed as $\min_{j \in \mathcal{R}_i}\{R_{ij}\}$. Following the same line of reasoning used in the proof of Theorem 2, the above expression has a finite expectation as long as $\sum_{j \in \mathcal{R}_i} \alpha_{ij} > 1 + |\mathcal{R}_i|$ (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

Fig. 1. Social forwarding at a glance
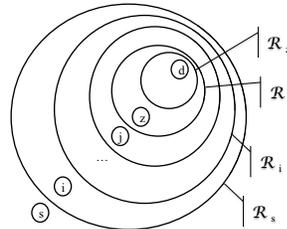
currently holding the message can only deliver it to the destination directly. Thus, $\alpha_{jd} > 2$ is required after $n - 1$ relays have been reached. Let us split all possible relays in $\mathcal{R}_s$ into two subsets $\mathcal{C}$ and $\mathcal{D}$, such that $\mathcal{C} \cup \mathcal{D} = \mathcal{R}_s$. 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. $\square$

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.

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. Both proofs 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 the proofs. Please note that in this case sequence π_i only contains nodes that belong to $\mathcal{R}_i$.

Lemma 4 (max_i^{sa}): When the social-aware multi-copy 2-hop forwarding protocol is in use and 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 z-th 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}): Under m-copy social-aware two-hop forwarding, 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 z-th 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.

Theorem 8 (m-copy 2-hop social-aware scheme): When the social-aware m-copy 2-hop forwarding protocol is used, the expected delay for messages generated by the source node s for the destination node d, with $m < |\mathcal{R}_s|$, achieves convergence if and only if the following condition holds true:

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

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

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.

Theorem 9 (m-copy n-hop social-aware scheme): When the multi-copy social-aware n-hop forwarding strategy is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if conditions C6 and C7 in Theorem 7 hold true.

Proof: The proof follows the same line of reasoning of the proof for Theorem 5 with the only difference that relays are selected in $\mathcal{R}_s$. Thus, we omit it. $\square$

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 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, in Lemma 6 and Theorem 10 below (whose proofs can be found in Appendix D) we tackle this problem from an analytical perspective, and we find that there is no winner in this case, and both the social-aware m-copy 2-hop scheme and the social-oblivious one can achieve convergence when the other one does not. In Appendix E we provide a concrete example for both cases.

Lemma 6 (Comparison of min_i and max_i): The following relationship holds between min_i and max_i for the social-oblivious and the social-aware m-copy 2-hop schemes under any given node permutation π_i:

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

Lemma 6 tells us that, for a fixed sequence π_i of encounters at the source node, the maximum number of first-hop convergent copies that the source can send in the social-oblivious case is always greater than those that it can send in the social-aware case. Unfortunately, the situation is the same at the second hop: the minimum number of copies required for convergence at the second hop is always higher in the social-oblivious case, thus nullifying the advantage at the first hop. In the following theorem, expanding on the conditions in Lemma 6 we analytically express the tie between social-aware and social-oblivious m-copy 2-hop forwarding strategies.

Theorem 10: Since both the following configurations are feasible under the conditions in Lemma 6, 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 7), or vice versa (Equation 8), depending on the underlying mobility process.

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

$$min_i^{so} > max_i^{so} \geq max_i^{sa} \geq min_i^{sa} \qquad (8)$$

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 Appendix E we also provide two concrete examples in which i) the social-aware m-copy 2-hop scheme achieves convergence while the social-oblivious m-copy 2-hop scheme does not, and ii) the social-oblivious m-copy 2-hop scheme achieves convergences and the social-aware m-copy 2-hop scheme does not.

The mobility scenario we consider is described by the exponent matrix in Figure 2. 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

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

Fig. 2. Exponent matrix

1-copy n-hop social-aware scheme achieves convergence in this scenario.

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

We now focus on the m-copy 2-hop scheme, recalling that the social-oblivious version and the social-aware 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 Figure 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 (Figure 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 5 copies with finite first-hop expected delay (while 7 are required). Hence, convergence cannot be achieved in this case.

8 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
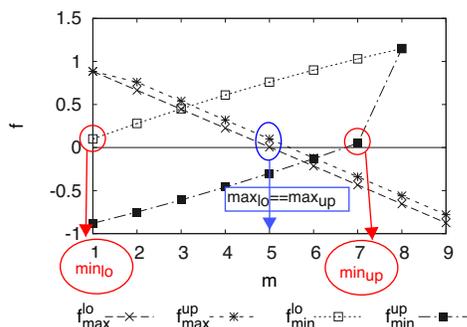
Fig. 3. min_i and max_i for the m-copy 2-hop scheme.

of forwarding protocols is finite. Our main result for the social-oblivious case is that convergence is not improved by an increased number of hops. Specifically, there is no advantage, as far as the convergence of the expected delay is concerned, in using more than two hops (and in some conditions direct transmission is the most efficient choice). In the social-aware case, instead, allowing more than two hops can provide convergence when all other strategies fail, because, when n is large enough, all nodes with a "bad" contact pattern with the destination can be bypassed.

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, 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 among the two, since either one can achieve convergence when the other fails depending on the underlying mobility scenario.

REFERENCES

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

[2] ——, "Efficient routing in intermittently connected mobile networks: The multiple-copy case," *IEEE/ACM Trans. on Netw.*, vol. 16, no. 1, pp. 77–90, 2008.

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

[4] C. Lee, "Heterogeneity in contact dynamics: helpful or harmful to forwarding algorithms in DTNs?" in *WiOPT'09*. IEEE, 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.*, pp. 1132–1147, 2009.

[6] A. Picu, T. Spyropoulos, and T. Hossmann, "An analysis of the information spreading delay in heterogeneous mobility dtns," in *IEEE WoWMoM*, 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. on Netw.*, vol. 14, no. 1, pp. 27–40, 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.*, pp. 606–620, 2007.

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

[10] W. Gao, Q. Li, B. Zhao, and G. Cao, "Multicasting in delay tolerant networks: a social network perspective," in *ACM MobiHoc'09*. ACM, 2009, pp. 299–308.

[11] 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," *Mobile Computing, IEEE Transactions on*, vol. 10, no. 12, pp. 1755–1768, 2011.

[12] 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. on Netw.*, vol. 17, no. 5, pp. 1578–1591, 2009.

[13] A. Passarella and M. Conti, "Analysis of individual pair and aggregate inter-contact times in heterogeneous opportunistic networks," *IEEE Trans. Mobile Comput.*, 2012. [Online]. Available: 10.1109/TMC.2012.213

[14] V. Conan, J. Leguay, and T. Friedman, "Characterizing pairwise inter-contact patterns in delay tolerant networks," in *Autonomics'07*, 2007.

[15] C. Lee and D. Eun, "Exploiting Heterogeneity in Mobile Opportunistic Networks: An Analytic Approach," in *IEEE SECON'10*. IEEE, 2010, pp. 1–9.

[16] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Tech. Rep., 2000.

[17] A. Lindgren, A. Doria, and O. Schelén, "Probabilistic routing in intermittently connected networks," *LNCS*, pp. 239–254, 2004.

[18] 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.

[19] E. Daly and M. Haahr, "Social network analysis for information flow in disconnected Delay-Tolerant MANETs," *IEEE Trans. Mobile Comput.*, pp. 606–621, 2008.

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

[21] 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, 2008.

[22] 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.

[23] 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, 2003.

[24] S. Ross, *Introduction to probability models*. Academic Press, 2007.

[25] M. Rytgaard, "Estimation in the Pareto distribution," *ASTIN Bulletin*, vol. 20, no. 2, pp. 201 – 216, 1990.

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

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

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

[29] A. Clauset, C. Shalizi, and M. Newman, "Power-law distributions in empirical data," *SIAM review*, vol. 51, no. 4, pp. 661–703, 2009.

[30] C. Boldrini, M. Conti, and A. Passarella, "Performance modelling of opportunistic forwarding under heterogenous mobility," IIT-CNR, Tech. Rep. TR-12/2013, http://cnd.iit.cnr.it/chiara/pub/techrep/boldrini2013heterogenous_tr.pdf.

Supplemental Material

APPENDIX A
MATHEMATICS OF POWER LAWS

In this appendix we summarize the properties of power law random variables that are used throughout the paper. Please recall that we consider Pareto random variables X_i whose CCDF is given by $F_{X_i}(x) = (\frac{b_i}{b_i+x})^{\alpha_i}$. Parameter b_i is referred to as *scale*, while α_i as *shape*. For ease of computation, when not stated otherwise, here we restrict to the case of power law random variables having the same scale b (i.e., $b_i = b, \forall i$).

Lemma A1 (Change of units for Pareto r.v.): Let us assume that c is a constant real value greater than zero and that X is a Pareto random variable with shape α and scale b. The distribution of $Y = X + c$ is equal to that of X shifted to the right by c. Thus, $F_Y(x) = F_X(x - c)$, or equivalently:

$$F_Y(x) = \left(\frac{b}{b + x - c}\right)^{\alpha} \qquad \text{(A.1)}$$

Proof: See [1] p.106. $\square$

Remark A1: A change of unit does not affect the conditions on the shape value α stated in Remark 1.

Lemma A2 (Conditioning on Pareto r.v.): Let us assume that c is a constant real value greater than zero and that X is a Pareto random variable with shape α and scale b. The distribution of X conditioned on the fact that $X > c$ is Pareto with shape α, scale $b + c$, and shift c.

Proof: We want to compute $P(X > x | X > c)$. By the rules of conditioning, we have that $P(A|B) = \frac{P(A \cap B)}{P(B)}$. This implies, in our case, $P(X > x | X > c) = 0$ for $x < c$, while $P(X > x | X > c) = \frac{P(X > y + c)}{P(X > c)}$ with $y = x - c$ (thus $y > 0$) for $x > c$. Thus, we get the following:

$$\frac{P(X > y + c)}{P(X > c)} = \left(\frac{b+c}{b}\right)^{\alpha} \left(\frac{b}{b+y+c}\right)^{\alpha} = \left(\frac{b+c}{y+b+c}\right)^{\alpha} \tag{A.2}$$

If we substitute back $y = x - c$, we obtain:

$$P(X > x | X > c) = \left(\frac{b+c}{(x-c)+b+c}\right)^{\alpha}, \qquad \text{(A.3)}$$

which is equivalent to a Pareto random variable with the same shape as X, scale $b + c$, and shift c. $\square$

Remark A2: As long as $c < \infty$, the expectation of a Pareto random variable X conditioned to be greater than c is finite under the same condition on the exponent value stated in Remark 1.

Figure A1 highlights the difference between a change of unit and conditioning. More specifically, X and $X + c$ are pretty much the same as long as x values are significantly greater than c. On the other hand, the conditioning operates a shift of the whole distribution.

Lemma A3 (Comparison between two Pareto r.v.): Let us consider two random variables, X_1 and X_2, following a power law distribution with shape α_1 and α_2, respectively, and same scale b. Then, the probability that X_1 is lower than X_2 is given by:

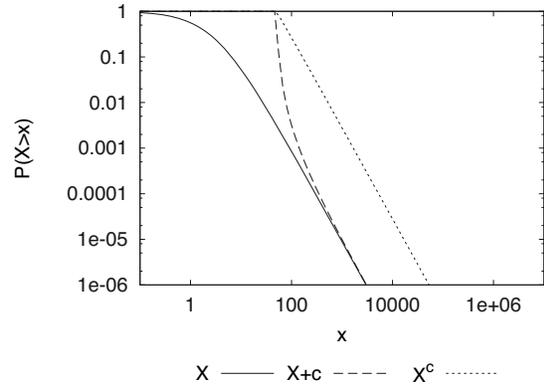$$P(X_1 < X_2) = \frac{\alpha_1}{\alpha_1 + \alpha_2} \qquad \text{(A.4)}$$

Fig. A1. Change of units vs conditioning ($\alpha = 2, b = 3, c = 50$)

Proof: We can rewrite $P(X_1 < X_2)$ using the law of total probability:

$$P(X_1 < X_2) = \int_0^{+\infty} P(X_1 < X_2 | X_2 = y) P\{X_2 = y\} dy$$
$$= \int_b^{+\infty} P(X_1 < y) P(X_2 = y) dy \qquad \text{(A.5)}$$

Equation A.4 is the solution to the above integral, computed after substituting the PDF and the CDF of the power law random variables into Equation A.5. $\square$

Corollary A1: Consider two random variables, X_1 and X_2, following a power law distribution with shape α_1 and α_2, respectively, scale b_1 and b_2, and shift c_1 and c_2. Then, the probability that X_1 is lower than X_2 (or vice versa) is always strictly greater than zero.

Proof: We can rewrite A.5 as $P(X_1 < X_2) = \int_0^\infty (1 - F_{X1}(x)) f_{X_2}(x) dx$, denoting with F_X and f_X the CCDF and PDF of X. Noting that X_1 cannot take values smaller than c_1 and that X_2 cannot take values smaller than c_2, we can rewrite the above equation as follows:

$$P(X_1 < X_2) = \int_{\max(c_1,c_2)}^\infty (1 - F_{X1}(x)) f_{X_2}(x) dx, \quad \text{(A.6)}$$

Defining $g(x) = (1 - F_{X1}(x)) f_{X_2}(x)$, the integral in Equation A.6 can be equal to zero only if $g(x) = 0$ for every x in $[\max(c_1, c_2), \infty)$. Clearly, this is not the case when X_1 and X_2 are Pareto random variables. In this case, in fact, function $g(x) = (1 - F_{X1}(x)) f_{X_2}(x)$ is positive and continuous in $(\max(c_1, c_2), \infty)$, because both the American Pareto CCDF and PDF are by definition positive and continuous in $(\max(c_1, c_2), \infty)$. Thus, $P(X_1 < X_2)$ is always greater than zero. If $P(X_1 < X_2) > 0$, then also $P(X_2 < X_1)$, which is equal to $1 - P(X_1 < X_2)$, is greater than zero. $\square$

A.1 Properties of the minimum of Pareto random variables

Lemma A4 (Minimum of n Pareto r.v. with the same scale): Random variable X defined as $X = \min_i\{X_i\}$, where

random variables X_i follow a power law distribution with shape α_i and scale b, is distributed according to a power law distribution with shape $\sum_i \alpha_i$ and scale b.

Proof: From standard probability theory we know that the CCDF of $\min_i\{X_i\}$ is equal to $\prod_i F_{X_i}$. When multiplying the CCDF of n power law random variables having the same scale, we again obtain a power law with shape equal to the sum of the shapes of the n power law random variables, and same scale. $\square$

Remark A3: The minimum of n random variables X_i each with shape α_i and scale b has a finite expectation as long as $\sum_i \alpha_i > 1$ (Remark 1).

Lemma A5 (Minimum of Pareto r.v. shifted by c): The minimum of n Pareto random variables shifted by the same value c is distributed as the minimum, shifted by c, of the n random variables. In other words, random variable X' defined as $\min\{X_i + c\}$, with c finite, is distributed as $X + c$, with $X = \min_i\{X_i\}$.

Proof: From Lemma A1, we know that $F_{X_i + c}(x) = \left(\frac{b}{b+x-c}\right)^{\alpha_i}$. Then $\prod_i F_{X_i+c} = \left(\frac{b}{b+x-c}\right)^{\sum_i \alpha_i}$, which corresponds to the CCDF of $X + c$. $\square$

Corollary A2 (Convergence of the min of r.v. shifted by c_i): The minimum $\min_i\{X_i + c_i\}$ of n Pareto random variables X_i each shifted by c_i shares the same convergence conditions for its expectation as the minimum $\min_i\{X_i\}$ of the same random variables unshifted.

Proof: By definition, inequality $X_i < c_i + X_i \leq c + X_i$ holds true, where $c = \max_i c_i$. From the above inequality we derive the following:

$$\min\{X_i\}_{i=1,\ldots,n} \quad < \quad \min\{c_i + X_i\}_{i=1,\ldots,n}$$
$$\min\{t_i + X_i\}_{i=1,\ldots,n} \quad \leq \quad \min\{c + X_i\}_{i=1,\ldots,n}$$

Applying, in order, Lemma A5, Remarks A1 and A3, we have that both $\min\{X_i\}_{i=1,\ldots,n}$ and $\min\{c + X_i\}_{i=1,\ldots,m}$ require $\sum_{i=1}^n \alpha_i > 1$ for convergence. Being $\min\{c_i + X_i\}_{i=1,\ldots,n}$ constrained by the two, we obtain that $\sum_{i=1}^n \alpha_i > 1$ is also a sufficient and necessary condition for the expectation of $\min\{c_i + X_i\}_{i=1,\ldots,n}$ to be defined. $\square$

Lemma A6 (Minimum of Pareto r.v. conditioned by c): The minimum of n Pareto random variables all conditioned to be greater than c is distributed as the minimum, conditioned to be greater than c, of the n random variables. In other words, random variable X' defined as $\min\{X_i^c\}$, with c finite, is distributed as X^c, with $X = \min_i\{X_i\}$.

Proof: From Lemma A2, we know that $F_{X_i^c}(x) = \left(\frac{b+c}{b+c+(x-c)}\right)^{\alpha_i}$. Then $\prod_i F_{X_i^c} = \left(\frac{b+c}{b+c+(x-c)}\right)^{\sum_i \alpha_i}$, which corresponds to the CCDF of X^c. $\square$

Corollary A3 (Converg. of the min of r.v. conditioned by c_i): Consider n Pareto random variables X_i, each with shape α_i and same scale b, conditioned to be greater than c_i (with $c_i \neq c_j$ if $i \neq j$ and $c_i < \infty, \forall i$). Then, the expectation of $\min_i\{X_i^{c_i}\}$ is defined under the same conditions for which the expectation of $\min_i\{X_i\}$ is defined, i.e., for $\sum_i \alpha_i > 1$.

Proof: For the sake of convenience and without loss of generality, in the following we assume that random variables are indexed by increasing value of c_i, i.e., $c_1 < c_2 < \ldots < c_n$. Exploiting Lemma C2, we can bound $\min_i\{X_i^{c_i}\}$ as follows:

$$\min_i\{X_i^{c_1}\} < \min_i\{X_i^{c_i}\} < \min_i\{X_i^{c_n}\}$$

From Lemma A6, we know that $\min_i\{X_i^{c_1}\}$ features a Pareto distribution with shape $\sum_i \alpha_i$, scale $b + c_1$, and shift c_1, whose expectation is defined as long as $\sum_i \alpha_i > 1$. Similarly, $\min_i\{X_i^{c_n}\}$ features a Pareto distribution with shape $\sum_i \alpha_i$, scale $b+c_n$, and shift c_n, whose associated condition is again $\sum_i \alpha_i > 1$. Being $\min_i\{X_i^{c_i}\}$ constrained by the two random variables, when they both converge $\min_i\{X_i^{c_i}\}$ must also converge (and vice versa). This implies that $\sum_i \alpha_i > 1$ is also the condition (necessary and sufficient) under which the expectation of $\min_i\{X_i^{c_i}\}$ is defined. $\square$

Remark A4 (Convergence of min of Pareto r.v.): The Pareto distribution resulting from the minimum of n Pareto distributions, each with its own shape α_i, is defined for $\sum_i \alpha_i > 0$ (due to the PDF normalization), and its expectation is finite when $\sum_i \alpha_i > 1$. The above conditions hold regardless of shifting and conditioning in the sense of Lemmas A5 and A6 and Corollaries A2 and A3.

APPENDIX B
PROOFS FOR SECTION 3

Lemma 1: 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 $R_{ij}^{t_i - t_0}$ (R_{ij} conditioned to be greater than $t_i - t_0$) or, equivalently from a convergence standpoint, as R_{ij}.

Proof: Let us assume that relay i has received the message at time t_i. The time before node i hands over the message to another potential relay j actually depends on whether i and j have met in time interval (t_0, t_i). If they did meet, say at $t^* \in (t_0, t_i)$, event happening with probability $P(R_{ij} < t_i)$, the time before i and j meet again is given by the intermeeting time between i and j conditioned to be greater than $t_i - t^*$ (we denote this quantity as $M_{ij}^{t_i - t^*}$). If they did not meet (event happening with probability $P(R_{ij} > t_i)$), we should consider the residual intermeeting time conditioned to be greater than $t_i - t_0$ ($R_{ij}^{t_i - t^0}$). This is due to the fact that from t_0 to the time of the first encounter the message sees the network as a random observer would see the renewal process describing encounters between i and j (hence residual intermeeting times are considered). Instead, after the first encounter, the message is not anymore a random observer, since a renewal has taken place, so the time between encounters is measured in terms of intermeeting times rather than residual intermeeting times. The two cases are summarized in Figure B2.
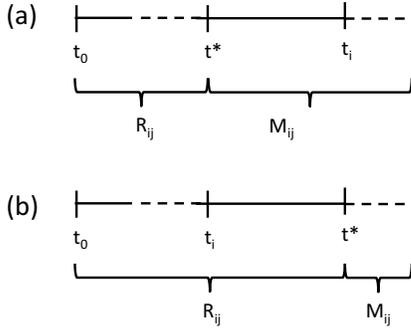
Fig. B2. Example of interplay between intermeeting times and residual intermeeting times when (a) $t^* \in (t_0, t_i)$ (b) $t^* > t_i$

Both $M_{ij}^{t_i - t^*}$ and $R_{ij}^{t_i - t^0}$ are Pareto distributed and their convergence is not affected by conditioning (Remark A2). However, convergence is more difficult for $R_{ij}^{t_i - t^0}$ (Remark 2). Since both cases are possible (in fact, $P(R_{ij} > t_i) > 0$ and $P(R_{ij} < t_i) > 0$, when R_{ij} is Pareto and t_i is finite), from a convergence standpoint we can restrict the analysis to the residual intermeeting time. In fact, if convergence is achieved when considering residual intermeeting times, convergence is guarantee also in the best case in which nodes i and j have met in (t_0, t_i) and the simple intermeeting time should have been considered instead. $\quad\square$

APPENDIX C
PROOFS FOR SECTION 5

We first prove Lemmas C1 and C2 that will be used throughout this section.

Lemma C1 (Set cardinality and minimum): Let us define random variable $X_{(n)}$ as the minimum of a variable number n of independent random variables, i.e, $X_{(n)} = \min_{i=\{1,\dots,n\}}\{X_i\}$. Regardless of the distribution of X_i, it holds that $X_{(n)} \geq X_{(n+1)}$, i.e., $P(X_{(n)} > x) \geq P(X_{(n+1)} > x), \forall n \geq 1$. In addition, if X_i can only take non-negative values, then $E[X_{(n)}] \geq E[X_{(n+1)}], \forall n \geq 1$.

Proof: $P(X_{(n)} > x) \geq P(X_{(n+1)} > x)$ follows from the fact that probabilities are constrained within zero and one, whatever the X_i. In fact, $P(X_{(n)} > x) = \prod_{i=\{1,\dots,n\}} P(X_i > x)$ and $P(X_{(n+1)} > x) = \prod_{i=\{1,\dots,n+1\}} P(X_i > x)$. Thus, $P(X_{(n)} > x) \geq P(X_{(n+1)} > x)$ is equivalent to $\prod_{i=\{1,\dots,n\}} P(X_i > x) \geq \prod_{i=\{1,\dots,n+1\}} P(X_i > x)$, or, alternatively:

$$\prod_{i=\{1,\dots,n\}} P(X_i > x) \geq P(X_{n+1} > x) \prod_{i=\{1,\dots,n\}} P(X_i > x).$$

The latter is always verified because $P(X_{n+1} > x) \in [0,1]$ by definition. In summary, $P(X_{(n)} > x)$ decreases as we increase n, i.e., as the cardinality of the set of random variables considered increases. For X_i taking only non-negative values, the expectation $E[X_{(n)}]$ can be computed as $\int_0^\infty P(X_{(n)} > x)dx$. Exploiting the fact that integration preserves the ordering of functions (i.e., if

$f(x) \geq g(x)$ in a generic interval $[a,b]$, then $\int_a^b f(x) \geq \int_a^b g(x))$, we also obtain that $E[X_{(n)}] \geq E[X_{(n+1)}], \forall n \geq 1$. $\quad\square$

Lemma C2 (Ordering M and R): Consider random variable M, featuring a power law distribution with shape α and scale t_0, and its residual R. Denote with M^{t^*} and R^{t^*} the random variables obtained when conditioning, respectively, M and R to be greater than a generic t^*. Then, for all $\alpha > 1$ and $t_0 > 0$, the following stochastic ordering applies:

$$M \leq M^{t_1} \leq M^{t_2} \leq R^{t_2} \leq R^{t_3}, \text{where } t_0 < t_1 < t_2 < t_3.$$

Proof: Recall that, given two random variables X and Y, $X \leq Y$ if $P(X > x) \leq P(Y > x)$ [2]. Let us first focus on $M \leq M^{t_1}$, with $t_1 > t_0$. Due to Lemma A2, M^{t_1} is again power law distributed, with shape α, scale $t_0 + t_1$, and shift t_1. When comparing the CCDF of M and M^{t_1} we obtain the following:

$$\begin{cases} \left(\frac{t_0}{t_0+t}\right)^\alpha \leq 1 & 0 \leq t < t_1 \\ \left(\frac{t_0}{t_0+t}\right)^\alpha \leq \left(\frac{t_0+t_1}{t_0+t}\right)^\alpha & t \geq t_1 \end{cases}$$

While the first inequality is trivial, the second one follows from the fact that function $f(t^*) = \frac{t_0+t^*}{t_0+t}$ is monotonically increasing with t^* when $t > 0$, and exponentiation preserves such property as long as $\alpha > 0$ (recall that in our case $\alpha > 1$). Applying the same reasoning, we also obtain $M^{t_1} \leq M^{t_2}$. Next, we compare M^{t_2} and R^{t_2}, the latter being Pareto distributed with exponent $\alpha - 1$, scale $t_0 + t_2$, and shift t_2 (Lemma A2). For $t \geq t_2$, we have $\left(\frac{t_0+t_2}{t_0+t}\right)^\alpha \leq \left(\frac{t_0+t_2}{t_0+t}\right)^{(\alpha-1)}$, because $\frac{t_0+t_2}{t_0+t}$ belongs to the interval $[0,1]$. Finally, using again the same approach, we have $\left(\frac{t_0+t_2}{t_0+t}\right)^{(\alpha-1)} \leq \left(\frac{t_0+t_3}{t_0+t}\right)^{(\alpha-1)}$, for all $t_3 > t_2$. $\quad\square$

C.1 Single-copy

Theorem 3 (1-copy n-hop scheme): When the single-copy n-hop relaying protocol is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if conditions C1 and C2 in Theorem 2 hold true.

Proof: The proof is composed of four parts. We first study the delivery from the source node to the relays (Part 1), then we concentrate on the delivery from the last relay to the destination node (Part 2), and after that we study the delivery from relay to relay along the multi-hop path (Part 3). Finally, in Part 4, we prove that any node $i \notin \{s,d\}$ has a non negligible probability of being the k-th hop along the n-th hop path, with $k \in \{1, \dots, n-1\}$, result which is at the basis of the derivations in the previous parts of the proof.

1) From source to relay. See proof of Theorem 2. ∎

2) From relay to destination. Let us now focus on the delivery from the last relay to the destination node. The time before the message is handed over by the last relay

(a generic node j) to node d can be modeled (Lemma 1) as R_{jd}. Thus, from Remark 2 condition $\alpha_{jd} > 2$ follows. We show in Part 3 of the proof that all nodes have a non negligible probability of being the $(n-1)$-th hop, thus condition $\alpha_{jd} > 2$ must be satisfied for all nodes $j \in \mathcal{P}_s - \{d\}$ (condition C2). ∎

3) From relay to relay. Now we discuss how to model the forwarding on intermediate relays. Assume that at time t_k a generic node i has just received the message from the $(k-1)$-th relay, thus becoming the k-th relay for the message. We denote with $\mathcal{E}_i^{t_k}$ the set of nodes that have been encountered by node i in the interval (t_0, t_k) ($\mathcal{E}_i^{t_k} \subseteq \mathcal{P}_i$).

In the memoryful case a relay can be used at most once. Thus, the k-th relay (assume it to be a generic node i) can only select as next hop one of the nodes that have not been used so far. We denote with $\mathcal{P}_i^k$ the set of nodes still available as relays when k hops have been exploited. The time $T_{i(k)}$ before node i, as the k-th relay, hands over the message to one of the available relays can be modeled, exploiting Lemma 1, as follows:

$$T_{i(k)} = \min\{R_{ij}\}_{j \in \mathcal{P}_i^k},$$

since node i will hand over the message to the first encountered node in $\mathcal{P}_i^k$. In the above equation, the cardinality of the set from which we take the minimum decreases as k increases, because, since relays cannot be used twice, at the k-th hop there are less relays available than at any previous hop. From Lemma C1 we know that convergence conditions are worse when the cardinality of the set from which we take the minimum is smaller. The worst case is thus for $k = n-1$, because in this case the only relay available is the destination. If we ensure convergence in this case, convergence in all other cases automatically follows. We derived in Part 2 of the proof that the convergence condition for this case is $\alpha_{id} > 2$ for all nodes $i \in \mathcal{P} - \{d\}$. ∎

4) Non negligible selection probability. In this part of the proof we want to show that all nodes but the source node have a non negligible probability of being selected as the $(k+1)$-th relay. To this aim, instead of exploiting Lemma 1, we model accurately $T_{i(k)}$ before node i, as the k-th relay, hands over the message to one of the available relays. Let us assume that the k-th relay i has received the message at time t_k. If node i met node j in (t_0, t_k), then the time before the two nodes meet again can be modeled as the intermeeting time M_{ij} conditioned on the time $t_{last(i,j)}$ since the last meeting between i and j. Thus, recalling that we denote with $\mathcal{P}_i^k$ the set of nodes still available as relays when k hops have been exploited and with $\mathcal{E}_i^{t_k}$ the set of nodes that have been encountered by node i in the interval (t_0, t_k), we can model $T_{i(k)}$ as follows:

$$T_{i(k)} = \min\{\{R_{ij}^{t_k - t_0}\}_{j \in \mathcal{P}_i^k - \mathcal{E}_i^{t_k}}, \{M_{ij}^{t_{last(i,j)}}\}_{j \in \mathcal{P}_i^k \cap \mathcal{E}_i^{t_k}}\}. \tag{C.1}$$

Now consider a specific node z. The probability p_{iz} that z is selected as next hop by current relay i can

be obtained from Equation C.1 as explained below. We define random variable Y_{iz} as follows:

$$Y_{iz} = \begin{cases} M_{iz}^{t_{last(i,z)}} & \text{if } z \in \mathcal{E}_i^{t_k} \\ R_{iz}^{t_k - t_0} & \text{otherwise.} \end{cases}$$

Y_{iz} describes the time before nodes i and z meet again. Please note that Y_{iz} is a Pareto random variable. Node z is selected as next hop with the following probability:

$$\begin{aligned} p_{iz} &= P\Big(Y_{iz} < \min\{\{R_{ij}^{t_k - t_0}\}_{j \in \mathcal{P}_i^k - \mathcal{E}_i^{t_k} - \{z\}}, \\ &\qquad \{M_{ij}^{t_{last(i,j)}}\}_{j \in \mathcal{P}_i^k \cap \mathcal{E}_i^{t_k} - \{z\}}\}\Big), \end{aligned} \tag{C.2}$$

where the righthand side of the inequality corresponds to the $T_{i(k)}$ in Equation C.1 from which we have removed the term corresponding to node z. We want to prove that the above probability is greater than zero. To this aim, please note that, according to Lemma C2, the righthand side of the inequality in Equation C.2 can be lower bounded by random variable X_i^*, defined as follows:

$$X_i^* = \min\{\{R_{ij}^{t_i^*}\}_{j \in \mathcal{P}_i^k - \mathcal{E}_i^{t_k} - \{z\}}, \{M_{ij}^{t_i^*}\}_{j \in \mathcal{P}_i^k \cap \mathcal{E}_i^{t_k} - \{z\}}\},$$

where $t_i^* = \min_{j \in \mathcal{P}_i^k \cap \mathcal{E}_i^{t_k} - \{z\}}\{t_{last(i,j)}\}$. From Lemma A4 we know that X_i^* features a Pareto distribution with scale $t_i^* + t_{min}$, shift t_i^*, and shape $\alpha_i^* = \sum_{j \in \mathcal{P}_i^k - \mathcal{E}_i^{t_k} - \{z\}} (\alpha_{ij} - 1) + \sum_{j \in \mathcal{P}_i^k \cap \mathcal{E}_i^{t_k} - \{z\}} \alpha_{ij}$. Then, Corollary A1 tells us that $P(Y_{iz} < X_i^*)$ is always greater than zero. Thus, also p_{iz} in Equation C.2 is greater than zero. This proves that all nodes have a non negligible probability of being selected as the $(k+1)$-th relay, with $k = 0, 1, \ldots, n-2$. ∎

Summarizing, we have derived the convergence conditions for the source-to-relay delivery (Part 1), for the relay-destination delivery (Part 2), and for the relay-to-relay delivery (Part 3). Of these conditions, those required for the the source-to-relay delivery ($\sum_{j \in \mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$) and for the relay-destination delivery ($\alpha_{id} > 2$ for all nodes $i \in \mathcal{P} - \{d\}$) automatically guarantee that the relay-to-relay delivery also converges. Thus, overall we require the same conditions C1 and C2 that we had derived for the single-copy 2-hop case.

Please note that these conditions are not only sufficient but also necessary. In fact, if condition C1 is not satisfied, the message cannot leave the source node within a finite expected time. Similarly, if condition C2 is not satisfied for all nodes $i \in \mathcal{P} - \{d\}$ there is a non negligible probability that the message reaches a node that, as last relay, is not able to reach the destination with a finite expectation. When this happens, the overall expected delay will diverge. □

C.2 Two-hop multi-copy schemes

Lemma 2 (max_i^{so}): When the multi-copy social-oblivious two-hop forwarding protocol is in use and 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:

$$max_i^{so} = \arg\min_m \{f_{max}^{so}(m, \pi_i) > 0\}, \qquad (2)$$

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 z-th node belonging to π_i.

Proof: According to the memoryful multi-copy two-hop relaying protocol, 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 k-th copy, t_0 the time at which the message is generated at the source, and t_k the time at which the k-th 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 k-th copy is relayed is given by $\min_{j \in \mathcal{P}_s^k} \{R_{sj}\}$. Thus, from Remark A4, we know that convergence is ensured 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 tells us that the smaller the cardinality of the set of random variables of which we take the minimum, the slower the convergence. This implies that the strictest condition for the convergence of the expected delay of the first hop is imposed by the m-th copy, i.e., by the one that sees that narrower set of nodes left for relaying. Thus, if we are able to define a convergence condition for the m-th 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 m-th copy. When the $(m - 1)$-th copy has been delivered, there are $N - 1 - (m - 1) = N - m$ potential relays left for the m-th 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[1] 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 i-th 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

1. 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.

now define a mapping $g^{(i)}$ that goes from set $\{\alpha_{sj}\}_{j \in \mathcal{P}_s}$ to set $\{\alpha_z^{(i)}\}_{z \in \{1, \dots, |\mathcal{P}_s|\}}$, where $\alpha_z^{(i)}$ corresponds to the exponent α_{sj} of the z-th element in π_i. Using the above notation, the time before the m-th copy is handed over is described by $\min_{j \in \nu_i} R_{sj}$. Using Remark A4 and the mapping defined above, we have that the convergence condition for the expected delay of the m-th copy is given by the following:

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

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

The above equation defines the convergence condition for the m-th 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 Lemma C1, convergence becomes more difficult as m increases. This is highlighted also by Equation C.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 2 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 C.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: In the previous proof we have studied what is the maximum number (max_i^{so}) of convergent copies that the source node can send, *given a sequence of encounters π_i*. This implies that to each π_i corresponds a value max_i^{so}. Here we want to investigate what are the smallest and greatest values that max_i^{so} can take and what are the π_i for which these values are achieved.

Let us start with max_{lo}^{so}. We want to find the permutation π_{i*} such that function $f_{max}^{so}(m, \pi_{i*})$ is smaller than, or at most equal to, any other $f_{max}^{so}(m, \pi_i)$ with $i \in \{1, \ldots, (N-1)!\}$. It is straightforward to prove that this happens when π_{i*} is such that nodes in π_{i*} are ordered in decreasing order of their exponents (in other words, $\alpha_z^{(i*)} \geq \alpha_j^{(i*)}$ for all $z < j$). In fact, when relays are encountered (and thus exploited) in decreasing order of their exponents, the relays left to the next copy are always those with the smallest exponents, hence the summation in Equation C.2 takes its minimum value. Being $f_{max}^{so}(m, \pi_i)$ decreasing with m, $f_{max}^{so}(m, \pi_{i*}) \leq f_{max}^{so}(m, \pi_i)$ implies that $f_{max}^{so}(m, \pi_{i*})$ crosses the x-axis before any other $f_{max}^{so}(m, \pi_i)$, thus providing the lower bound on the maximum number of convergent first hop copies that the source node can send in the worst case.

The derivation for max_{up}^{so} is analogous. We want to find the permutation $\pi_{i'}$ such that function $f_{max}^{so}(m, \pi_{i'})$ is greater than, or at most equal to, any other $f_{max}^{so}(m, \pi_i)$ with $i \in \{1, \ldots, (N-1)!\}$. The summation in Equation C.2 takes its maximum value when exponents $\alpha_z^{(i)}$ ($z = m, \cdots, |\mathcal{P}_s|$) are the largest possible, thus when $\pi_{i'}$ corresponds to an increasing order of exponents α_{sj}. Then, as $f_{max}^{so}(m, \pi_{i'}) \geq f_{max}^{so}(m, \pi_i)$, $f_{max}^{so}(m, \pi_{i'})$ will cross the x-axis after any other $f_{max}^{so}(m, \pi_i)$. The last integer m value before or at the intersection will provide max_{up}^{so}. □

Lemma 3 (min_i^{so}): Under multi-copy social-oblivious two-hop forwarding, 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 z-th node in encounter sequence $\pi_i - \{d\}$.

Proof: In this lemma we focus on the delivery from intermediate relays to the destination. In the memoryful case, the number of intermediate relays in use is equal to the number of copies sent by the source nodes, and relays are distinct by definition. Moreover, for a given sequence π_i of encounters at the source node, the identities of relays are fixed. Specifically, nodes acting as relays are the first m nodes in π_i. Since we are focusing on the delivery from intermediate relays to the destination, we have to exclude the destination from the set of possible relays (it was however considered when studying first hop convergence in Lemma 2, so the contribution of the direct delivery from source to destination has been already taken into account), hence we are considering permutations π_i without the destination node.

Assume that we have the m relays holding a copy of the message. The distinct m relays have received the copy to relay at different times. More specifically, the first relay received its copy at a time t_1 that is smaller than time t_2 at which the second relay has received its copy, and so on. Thus, assuming again that t_0 is the message creation time and denoting with $R_{(z)}$ the residual inter-meeting time between the z-th relay (corresponding to the z-th node in π_i) and the destination (and with $\alpha_{(z)}$ its shape), we have that in each interval $(t_k, t_{k+1}]$ the time before the message is delivered to the destination by the k nodes currently holding a copy of the message is given, in the worst case (Lemma 1), by $\min\{R_{(1)}^{t_k}, ..., R_{(k)}^{t_k}\}$. Thus, after the m-th copy has been relayed, the expectation of the delay for the second hop is defined as long as the expectation of $\min\{R_{(1)}^{t_1}, ..., R_{(m)}^{t_m}\}$ is defined. We do not need to require convergence before t_m, time at which the last copy is relayed by the source because t_m is finite by definition, but we have to make sure of convergence after t_m. From Corollary A3 we obtain that convergence is ensured as long as the following holds:

$$\sum_{z=1}^m \alpha_z^{(i)} - m - 1 > 0. \qquad (C.3)$$

Analogously to what we did in the proof of Lemma 2, we use the above equation for deriving the minimum number of relays required for the convergence to be achieved under meeting sequence π_i. More specifically, let us denote the left-hand side of the equation as $f_{min}^{so}(m, \pi_i)$. It is easy to prove that $f_{min}^{so}(m, \pi_i)$ increases with m. In fact, consider moving from $f_{min}^{so}(m, \pi_i)$ to $f_{min}^{so}(m+1, \pi_i)$. The difference $f_{min}^{so}(m+1, \pi_i) - f_{min}^{so}(m, \pi_i)$ is equal to $\alpha_{m+1}^{(i)} - 1$, which is always greater than zero, since we assume that all exponents are bigger than 1. Given that $f_{min}^{so}(m, \pi_i)$ is increasing, either it starts above zero (and it will never go back) or it will start below, thus crossing the x axis at some point. Thus, the m value corresponding to the first integer for which $f_{min}^{so}(m, \pi_i) \geq 0$ gives min_i^{so}.

□

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.

Proof: The proof is similar to the proof for Corollary 1. Starting with the upper bound on min_i^{so} (corresponding to the worst case from the convergence standpoint), we look for the permutation π_{i*} such that $f_{min}^{so}(m, \pi_{i*})$ is smaller than any other $f_{min}^{so}(m, \pi_i)$. In this case, in fact, $f_{min}^{so}(m, \pi_i^*)$ crosses the x-axis after all others and thus yields the highest min_i^{so}. Since $f_{min}^{so}(m, \pi_i)$ is always increasing with m, this worst case happens when π_{i*} corresponds to nodes $j \in \mathcal{P}_s$ ordered with increasing exponents α_{jd}. In fact, the lower the first m exponents in π_i, the more difficult the convergence and the slower $f_{min}^{so}(m, \pi_i)$ to increase. Thus, the upper bound min_{up}^{so} is achieved when π_{i*} corresponds to nodes $j \in \mathcal{P}_s$ ordered with increasing exponents α_{jd}.

An analogous derivation applies to the best case, with the difference that this time we are interested in the permutation $\pi_{i'}$ such that $f_{min}^{so}(m, \pi_{i'})$ is greater than any other $f_{min}^{so}(m, \pi_i)$. In fact, in this case $f_{min}^{so}(m, \pi_{i'})$ crosses the x-axis before all others, thus yielding the smallest min_i^{so}. We know that $f_{min}^{so}(m, \pi_i)$ grows with m and it grows faster as the summation $\sum_{z=1}^{m} \alpha_z^{(i)}$ in it yields bigger results. The fastest growth for the summation is when it involves the largest exponents available. Thus, the lower bound min_{lo}^{so} is achieved when $\pi_{i'}$ corresponds to nodes $j \in \mathcal{P}_s$ ordered with decreasing exponents α_{jd}. $\square$

Theorem 4 (m-copy 2-hop scheme): When the m-copy two-hop forwarding protocol is used (with $m < N - 1$), the expected delay for messages generated by the source node s for the destination node d converges if and only if the following condition holds true:

C3 $\quad m \geq min_{up}^{so} \wedge max_i^{so} \geq min_i^{so}, \forall i \in \{1, \ldots, |\mathcal{P}_s^p|\}$, where set $\mathcal{P}_s^p$ is the set of all permutations for elements in $\mathcal{P}_s$.

Proof: The proof has been provided in the body of the paper. Here we just want to highlight that condition C3 is both necessary and sufficient. To this aim, we split it in two parts: $m \geq min_{up}^{so}$ (condition C3.a) and $max_i^{so} \geq min_i^{so}$ (condition C3.b). Condition C3.a is necessary since, if $m < min_{up}^{so}$, in some cases the source node will not *generate* enough copies for the second hop to converge. Since convergence must be ensured in all cases, m should be set to a value greater than the number of copies required at the second hop in the worst case. On a similar note, if condition C3.b were not true (i.e, $\exists i \in \{1, \ldots, |\mathcal{P}_s^p|\}$ such that $max_i^{so} < min_i^{so}$) it would imply the existence of a path in which the source node is not able to send the minimum number of *convergent* copies that guarantees convergence at the second hop. Since any path can be taken with non negligible probability, the existence of a divergent path implies the divergence of the overall expected delay. $\square$

C.2.1 Memoryful m-copy 2-hop forwarding with homogeneous encounters: a comparison with [3]

As discussed before, Chaintreau et al. [3] studied the m-copy memoryful two-hop scheme under homogeneous mobility patterns (corresponding to $\alpha_{ij} = \alpha, \forall i, j$). For the sake of completeness, in Corollary 3 we verify that Theorem 4 confirms and extends the results in [3].

Corollary 3: In a homogeneous network where the intermeeting times M_{ij} follow a power law distribution with shape α for all i, j node pairs, when the social-oblivious m-copy two-hop strategy is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if the following holds true:

$$\begin{cases} \alpha > \frac{1}{m} + 1, & m \leq \frac{N}{2} \\ \alpha > \frac{2}{N} + 1, & m > \frac{N}{2} \end{cases} . \qquad \text{(C.4)}$$

Proof: We start by noting that in a homogeneous network all relay permutations π_i determine the same
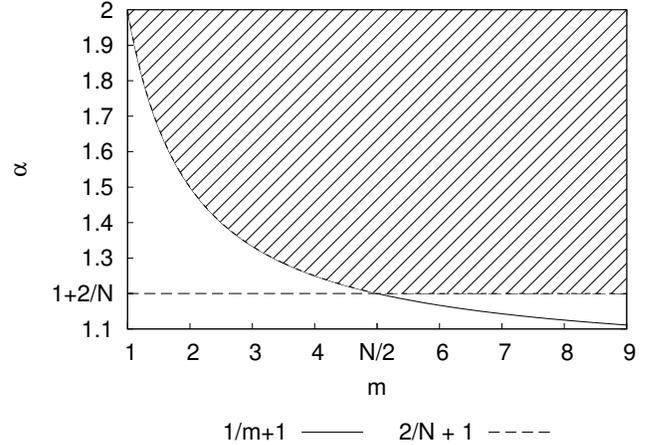

Fig. C3. Convergence for the m-copy 2-hop scheme in the homogeneous case (N=10)

ordering on $\{\alpha_{sj}\}_j$ and $\{\alpha_{jd}\}_j$, since exponents are all the same. Thus, we have $f_{max}^{so}(m, \pi_i) = f_{max}^{so}(m, \pi_j) = m + (N - m)\alpha - (N + 1), \forall i, j$. Similarly $f_{min}^{so}(m, \pi_i) = f_{min}^{so}(m, \pi_j) = m\alpha - (1 + m), \forall i, j$. Now, let us split condition C3 in two parts: $m \geq min_{up}^{so}$ (condition C3.a) and $max_i^{so} \geq min_i^{so}$ (condition C3.b). For condition C3.b to be satisfied, we need to impose condition $N - \frac{1}{\alpha-1} \geq \frac{1}{\alpha-1}$, from which we obtain $\alpha \geq \frac{2}{N} + 1$. For condition C3.a we need to require $m \geq \frac{1}{\alpha-1}$, from which we get $\alpha \geq 1/m + 1$. These two conditions must hold simultaneously for convergence to be achieved, from which Equation C.4 follows (it delimits the shadowed area corresponding to the stability region in Figure C3). $\square$

An important consequence of the above corollary is that if $\alpha \leq 1 + \frac{2}{N}$, the m-copy two-hop strategy will diverge regardless of the m value (Figure C3). Beyond $1 + \frac{2}{N}$, there always exists a m value for which convergence is achieved. Finally, please note that the necessary and sufficient condition in Equation C.4 confirms and extends the sufficient condition provided by Chaintreau et al. [3]. In fact, Chaintreau et al., under the assumption $N \geq 2m$ (corresponding to our $m \leq \frac{N}{2}$ case), derived that the expected delay of the m-copy two-hop scheme converges in a homogeneous setting as long as $\alpha > 1 + \frac{1}{m}$, which is in agreement with our result.

C.3 Multi-hop multi-copy schemes

Theorem 5: When the social-oblivious m-copy n-hop protocol is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if condition C1 and C2 in Theorem 2 hold true.

Proof: The proof consists in showing that there is a non negligible probability of path merging, i.e., there is a non negligible probability that all copies are handed over to the same relay at the second forwarding step. Path merging affects the delivery to the destination

because, as we have proved in Theorem 4, the more the distinct copies of the message, the better the convergence conditions. If we prove that path merging happens with non negligible probability, in the worst case the delivery to the destination must be treated as a single-copy n-hop delivery (hence conditions C1 and C2), i.e., we only have to guarantee that at least one copy is delivered with finite expected delay, because in the worst case one copy is all we have after the first hop. Before proceeding into the proof, please recall that each copy keeps track of used relays in order to guarantee that the same copy is not relayed twice by the same node, which leads to copies traveling along independent paths. However, it is perfectly possible that two copies are relayed to the same node, which is a new (not previously used) relay for both.

The first hop (i.e., the delivery from the source node to the first relay for each of the m copies) follows the same rules described in Lemma 2. Thus, in the following we focus on what happens after the m copies have been handed over to their first hop relays. More specifically, we prove that, at this first hop, there is already a non negligible probability of path merging. We know that the probability that a generic node z is the k-th relay is simply the probability that z is selected as next hop by the current relay i, which is, in turn, equivalent to the probability that z is the first node encountered by i not already used as relay for that copy. We denote with $\mathcal{P}_i^{k,c}$ the set of relays still available to node i after the c-th copy has travelled k hops. Since we focus on the delivery from the first relay to the second one, we consider $\mathcal{P}_i^{1,c}$. Thus, $z \in \mathcal{P}_i^{1,c}$.

Recall that, in the time interval between message generation at t_0 and the reception of the c-th copy at t_i, node i might have met some nodes. We denote with $\mathcal{E}_i^{t_i}$ the set of nodes that have been encountered by node i in the interval (t_0, t_i) ($\mathcal{E}_i^{t_i} \subseteq \mathcal{P}_i$). When node z does not belong to $\mathcal{E}_i^{t_i}$ (i.e., when node i did not encounter z in the interval (t_0, t_i)), the time before nodes i and z meet is described by $R_{iz}^{t_i - t_0}$. Otherwise, the time before the two nodes meet is given by $M_{iz}^{t_{last(i,z)}}$, where $t_{last(i,z)}$ denotes the time passed since the last meeting between i and z. We define random variable Y_{iz} (with $z \in \mathcal{P}_i^{1,c}$) as follows:

$$Y_{iz} = \begin{cases} M_{iz}^{t_{last(i,z)}} & \text{if } z \in \mathcal{E}_i^{t_i} \\ R_{iz}^{t_i - t_0} & \text{otherwise} \end{cases}. \qquad \text{(C.5)}$$

Y_{iz} describes the time before i and z meet again and is a Pareto random variable. Similarly to Part 3 of the proof of Theorem 3, we derive that node z is selected as next hop with the following probability:

$$p_{iz}^{1,c} = P\left(Y_{iz} < \min\{\{R_{ij}^{t_i - t_0}\}_{j \in \mathcal{D}_i - \{z\}}, \{M_{ij}^{t_{last(i,j)}}\}_{j \in \mathcal{I}_i - \{z\}}\}\right), \qquad \text{(C.6)}$$

where $\mathcal{D}_i = \mathcal{P}_i^{1,c} - \mathcal{E}_i^{t_i}$ and $\mathcal{I}_i = \mathcal{P}_i^{1,c} \cap \mathcal{E}_i^{t_i}$. We want to prove that the above probability is greater than zero. To this aim, please note that, according to Lemma C2,

the right-hand side of the inequality in Equation C.6 can be lower bounded by random variable X_i^*, defined as follows:

$$X_i^* = \min\{\{R_{ij}^{t_i^*}\}_{j \in \mathcal{D}_i - \{z\}}, \{M_{ij}^{t_i^*}\}_{j \in \mathcal{I}_i - \{z\}}\}, \qquad \text{(C.7)}$$

where $t_i^* = \min_{j \in \mathcal{I}_i - \{z\}}\{t_{last(i,j)}\}$. From Lemma A4 we know that X_i^* features a Pareto distribution with scale $t_i^* + t_{min}$, shift t_i^*, and shape $\alpha_i^* = \sum_{j \in \mathcal{D}_i - \{z\}} (\alpha_{ij} - 1) + \sum_{j \in \mathcal{I}_i - \{z\}} \alpha_{ij}$. Corollary A1 tells us that $P\left(Y_{iz} < X_i^*\right)$ is always greater than zero. Hence, also $p_{iz}^{1,c}$ is always greater than zero.

Let us now denote with $\mathcal{C}$ the set of nodes that are current relays for the m copies (hence, $|\mathcal{C}| = m$). In order to prove that there is a non negligible probability at the first hop that all copies are relayed to the same node, we need $p_{worst} > 0$, where p_{worst} is defined as follows:

$$p_{worst} = \prod_{i \in \mathcal{C}} p_{iz}^{1,c} \qquad \text{(C.8)}$$

Basically, p_{worst} is the joint probability that all relays will hand over their copy to node z at the next forwarding opportunity. Given that $p_{iz}^{1,c}$ is always greater than zero, then also p_{worst} will be always greater than zero. Since there is a non negligible probability of path merging, in the worst case the forwarding process ends up with just one single copy after the first hop. In order for this copy to achieve a convergent expected delay we need to require those conditions that apply to the 1-copy n-hop delivery, which are conditions C1 and C2. □

APPENDIX D
PROOFS FOR SECTION 7

Lemma 6 (Comparison of min_i and max_i): The following relationship holds between min_i and max_i for the social-oblivious and the social-aware m-copy 2-hop schemes under any given node permutation π_i:

$$max_i^{so} \geq max_i^{sa} \wedge min_i^{so} \geq min_i^{sa}. \qquad \text{(6)}$$

Proof: In this proof we want to derive ordering relationships between the social-oblivious and the social-aware min_i and max_i. To this aim, recall that the social-oblivious protocols select their relays in set $|\mathcal{P}_s|$, while social-aware schemes select relays in $|\mathcal{R}_s|$.

Let us start with max_i. Based on Lemma C1, first hop convergence is more difficult for the social-aware schemes, since the cardinality of the set from which they select relays is smaller. Thus, in the following we consider the extreme case in which the social-aware approach is more advantaged, which is when $\{\alpha_{sj}\}_{j \in \mathcal{R}_s}$ corresponds to the highest $\alpha_{sj} \in \mathcal{P}_s$. If we find that inequality $max_i^{so} \geq max_i^{sa}$ holds in this case, it will also hold in all other cases. In order to penalize as much as possible the social-oblivious scheme, we assume that $\alpha_{sj} = 1 + \epsilon, \forall j \in \mathcal{P}_s - \mathcal{R}_s$, and we let ϵ approach zero. Thus, $\{\alpha_{sj}\}_{j \in \mathcal{P}_s - \mathcal{R}_s}$ contains the lowest possible exponents. At this point, let us split the analysis in two. First, we consider the worst case for the first

hop convergence, i.e., when relays are selected so that nodes with the highest exponents are taken first. Since the highest exponents are the same in $\mathcal{P}_s$ and $\mathcal{R}_s$, the behavior up to $m = |\mathcal{R}_s|$ is exactly the same. After that, the social-aware scheme stops forwarding, since there are no relays left in $\mathcal{R}_s$. Vice versa, the social-oblivious scheme can still select relays in $\mathcal{P}_s - \mathcal{R}_s$. Thus, the social-oblivious scheme can at least send as many copies as its social-aware counterpart ($max_{lo}^{so} \geq max_{lo}^{sa}$). Let us now focus on the best case for the first hop convergence. This corresponds to relays with the lowest exponents being selected first. In the social-oblivious case, until the lowest ones run out the highest ones remain available, thus improving convergence. Vice versa, in the social-aware case, nodes with the highest exponents are used from the beginning, since those are the only ones available. Thus, it follows that $max_{up}^{so} \geq max_{up}^{sa}$. The same line of reasoning that we have applied to the best and worst case relay selection can be generalized to whatever node permutation π_i. More specifically, we can ignore permutations in which the lowest exponents ($=1+\epsilon$) are taken among the first $|\mathcal{R}_s|$, since this selection favors the social-oblivious approach (for which all the highest exponents remains available for the next copies, while they are consumed in the social-ware case). For all other permutations, since the first part ($|\mathcal{R}_s|$) overlaps, and the second one gives a chance to the social-oblivious approach, ordering $max_i^{so} \geq max_i^{sa}$ holds. Summarizing, even in the most favorable case for the social-aware schemes, for all permutations the social-oblivious approach is able to send a maximum number of first hop convergent copies that is equal to or greater than that of the social-aware scheme.

Let us now focus on min_i. We know that, by definition, set $\{\alpha_{jd}\}_{j \in \mathcal{R}_s - \{d\}}$ (hereafter referred to as set $\mathcal{A}$) contains the highest exponents in $\{\alpha_{jd}\}_{j \in \mathcal{P}_s - \{d\}}$ (hereafter referred to as set $\mathcal{B}$). Then, consider the worst case from the min_i standpoint, corresponding to relays with the smallest exponents being taken first. From the above definition we know that the sum of the m smallest exponents in $\mathcal{A}$ is greater than or equal to the sum of the smallest exponents in $\mathcal{B}$. Thus, if it exists, $min_{up}^{sa} \leq min_{up}^{so}$. Please note that min_{up}^{sa} can be at most equal to $|\mathcal{R}_s - \{d\}|$, while min_{up}^{so} can reach $|\mathcal{P}_s - \{d\}|$. Thus, min_{up}^{so} can exists even if min_{up}^{sa} does not. Let us now consider the best case for min_i, corresponding to relays with the highest α_{jd} exponents being taken first. Since the set of the highest $|\mathcal{R}_s - \{d\}|$ exponents in $\mathcal{A}$ and $\mathcal{B}$ overlap, we have that $min_{lo}^{sa} = min_{lo}^{so}$, when min_{lo}^{sa} is defined (i.e., when the function crosses the x-axis in Equation 5). Note that also min_{lo}^{so} is defined in a larger interval than min_{lo}^{sa}. If we extend this analysis to a generic permutation (not only those corresponding to the best and worst case), we obtain similar results. Again we consider only permutations whose first $|\mathcal{R}_s - \{d\}|$ nodes are those in $\mathcal{R}_s - \{d\}$, since when this is not true the curve for the social-oblivious approach is always below the social-aware one, corresponding to a higher min_i^{so}.

Considering only the above subset of permutations, we have that min_i^{so} and min_i^{sa} are overlapping, as long as min_i^{sa} is defined. Summarizing the above results, we have that min_i^{so} is in general greater than or equal to min_i^{sa} for the same permutation π_i. $\square$

Theorem 10: Since both the following configurations are feasible under the conditions in Lemma 6, 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 7), or vice versa (Equation 8), depending on the underlying mobility process.

$$max_i^{so} \geq min_i^{so} \geq min_i^{sa} > max_i^{sa} \tag{7}$$
$$min_i^{so} > max_i^{so} \geq max_i^{sa} \geq min_i^{sa} \tag{8}$$

Proof: Theorem 10 simply follows from merging the inequalities in Lemma 6 with the conditions for divergence in C3 and C8 ($min_i^{so} > max_i^{so}$ and $min_i^{sa} > max_i^{sa}$, respectively). Here we assume that we are always able to set m to be greater than min_{up}^{so} if min_{up}^{so} is defined (the same holds for min_{up}^{sa}). $\square$

Corollary D1 (The $\alpha_{ij} < 1$ ($\forall i, j$) case): When $\alpha_{ij} \leq 1$ for all i, j node pairs, none of the forwarding strategies studied in this paper, either social-oblivious or social-aware, is able to achieve a finite expected delay.

Proof: We provide the proof for the social-oblivious case. For the social-aware case, the line of reasoning is the same, after substituting $\mathcal{P}_s$ with $\mathcal{R}_s$. The time before the source node hands over the first copy of the message to the first encountered node is given by $\min_{j \in \mathcal{P}_s} R_{sj}$, which converges as long as $\sum \alpha_{sj} > 1 + |\mathcal{P}_s|$. When $\alpha_{ij} \leq 1$ for all node pairs, α_{ij} can be at most equal to $1 - \epsilon$, with $\epsilon \to 0$. If we substitute this expression in the convergence condition, we obtain $(N - 1)(1 - \epsilon) > N$, which is never satisfied. Thus, if not even a single copy can be sent a finite expected time, the overall expected delay will surely diverge. $\square$

APPENDIX E
ADDITIONAL EXAMPLES

In Section 7.1 we have provided a concrete example in which the 1-copy n-hop scheme is the only one achieving convergence. In this second example we first study a scenario in which the social-oblivious m-copy 2-hop scheme wins over its social-aware counterpart in terms of convergence. We consider a network with 14 nodes, which meet according to exponent matrix in Figure E4. We assume that node 1 is the source node and node 14 is the destination node. Please note that the exponent values are purposely artificial, chosen in order to highlight the peculiarities of the different scenarios. Note also that we only report exponent values for source-relay and relay-destination pairs, since we are considering two-hop protocols (hence, relay-to-relay delivery is not allowed).

With the social-oblivious m-copy 2-hop scheme, the source node can hand over the message to any node it meets. Vice versa, with the social-aware m-copy 2-hop scheme, the source node can only hand over the

(a) The social-oblivious case
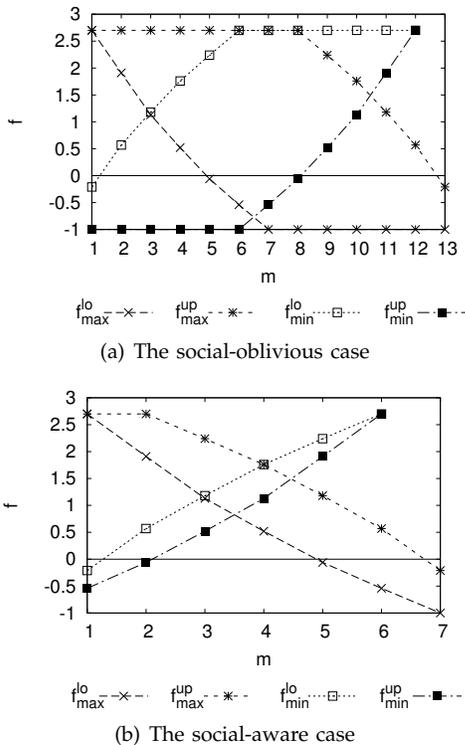

(b) The social-aware case

Fig. E5. Scenario 2, in which the social-aware approach reaches convergence while the social-oblivious one does not.

message to nodes with higher fitness (roughly corresponding to a higher Pareto exponent for meeting with the destination). In Figure E4, we have highlighted in blue the exponents α_{sj} corresponding to nodes j in set $\mathcal{R}_s$, while exponents α_{sj} corresponding to nodes j in set $\mathcal{P}_s$ simply corresponds to the first row of the matrix (minus the null element $\alpha_{1,1}$). As for the second hop, exponents α_{jd} associated with relays j in $\mathcal{R}_s$ have been highlighted with a dashed blue line, while exponents α_{jd} associated with relays j in $\mathcal{P}_s$ simply corresponds to the last column of the matrix (minus element $\alpha_{1,14}$ and null element $\alpha_{14,14}$). The main feature of this scenario is that nodes with the highest fitness (those in the dashed blue box) are also encountered frequently by the source, and, vice versa, nodes with low fitness are rarely encountered.

Let us now analyze the convergence in this scenario. As we have done in Section 7.1, 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. From Figure 5(b), we derive that $min_{up}^{sa} = 3$ and $max_{lo}^{sa} = 4$ in the social-aware case, thus sufficient condition C8$_{[s]}$ for convergence is satisfied. Vice versa, since $min_{up}^{so} = 9$ and $max_{lo}^{so} = 4$, sufficient condition C3$_{[s]}$ does not hold true. When computing the necessary and sufficient condition C3 we obtain that there exists a least one permutation $\pi_i = \{5, 12, 11, 8, 9, 10, 3, 4, 6, 7, 2, 13, 14\}$ such that $min_i^{so} > max_i^{so}$ (with $min_i^{so} = 7$ and $max_i^{so} = 6$), thus violating condition C3.

The net advantage of the social-aware approach in this

case is its ability to select only those nodes that have higher fitness and that are also frequently encountered by the destination. This allows the social-aware scheme to reach convergence when its social-oblivious counterpart does not.

The last scenario (Figure E6) we consider is one in which the social-oblivious m-copy 2-hop scheme performs better, from the convergence standpoint, than the social-aware one. The difference with respect to the previous case is that this time there is a negative correlation between the exponents characterizing meetings with the source node and the exponents characterizing meetings with the destination node.

In the social-oblivious case (Figure 7(a)), $min_{up}^{so} = 6$ and $max_{lo}^{so} = 6$, so the sufficient condition C3$_{[s]}$ is satisfied. In the social-aware case (Figure 7(b)), sufficient condition C8$_{[s]}$ is not satisfied, and neither the necessary and sufficient condition C8, since $min_i^{sa} = 1$ for all permutations but the source node is not even able to send just a single copy with finite first-hop delay.

(a) The social-oblivious case

(b) The social-aware case
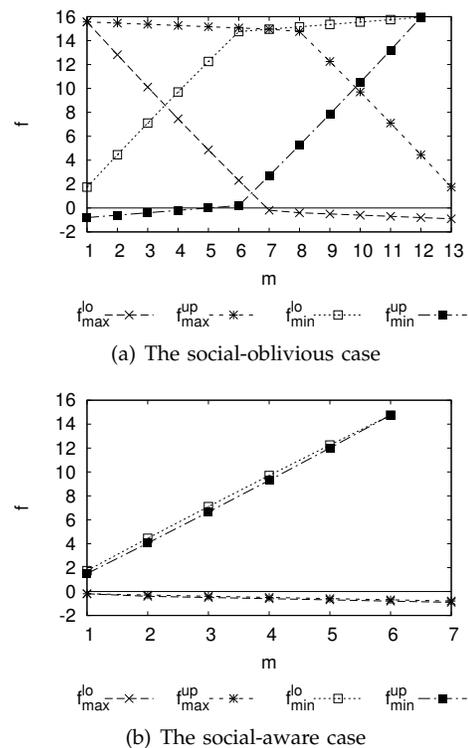
Fig. E7. Scenario 3, in which the social-oblivious approach reaches convergence while the social-aware one does not.

Contrary to the previous scenario, in this case the selection of nodes performed by the social-aware scheme is not effective. In fact, nodes with higher fitness are encountered only rarely by the destination, which makes extremely difficult for copies to leave the source.

REFERENCES

[1] M. Dekking, C. Kraaikamp, and H. Lopuhaa, *A modern introduction to probability and statistics: understanding why and how*. Springer Verlag, 2005.

$$
\begin{pmatrix}
- & 1.00001 & 1.00001 & 1.00001 & 1.00001 & 1.00001 & 1.00001 & \boxed{1.46 \ \ 1.48 \ \ 1.58 \ \ 1.61 \ \ 1.78 \ \ 1.79 \ \ 1.00002} \\
\cdots & - & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.0001 \\
\cdots & \cdots & - & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.0001 \\
\cdots & \cdots & \cdots & - & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.0001 \\
\cdots & \cdots & \cdots & \cdots & - & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.0001 \\
\cdots & \cdots & \cdots & \cdots & \cdots & - & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.0001 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & - & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.0001 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & - \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.46 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ - \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.48 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ - \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.58 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ - \ \ \cdots \ \ \cdots \ \ 1.61 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ - \ \ \cdots \ \ 1.78 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ - \ \ 1.79 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ -
\end{pmatrix}
$$

Fig. E4. Exponent matrix for which the social-aware approach reaches convergence while the social-oblivious one does not.

$$
\begin{pmatrix}
- & 3.5 & 3.55 & 3.6 & 3.65 & 3.7 & 3.75 & \boxed{1.1 \ \ 1.1 \ \ 1.1 \ \ 1.1 \ \ 1.1 \ \ 1.1 \ \ 1.21} \\
\cdots & - & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.2 \\
\cdots & \cdots & - & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.2 \\
\cdots & \cdots & \cdots & - & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.2 \\
\cdots & \cdots & \cdots & \cdots & - & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.2 \\
\cdots & \cdots & \cdots & \cdots & \cdots & - & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.2 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & - & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 1.2 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & - \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 3.5 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ - \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ 3.55 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ - \ \ \cdots \ \ \cdots \ \ \cdots \ \ 3.6 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ - \ \ \cdots \ \ \cdots \ \ 3.65 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ - \ \ \cdots \ \ 3.7 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ - \ \ 3.75 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ \cdots \ \ -
\end{pmatrix}
$$

Fig. E6. Exponent matrix for which the social-oblivious approach reaches convergence while the social-aware one does not.

[2] M. Shaked and J. G. Shanthikumar, *Stochastic orders*. Springer, 2007.
[3] 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.*, pp. 606–620, 2007.