Less is More: Long Paths do not Help the Convergence of Social-Oblivious Forwarding in Opportunistic Networks

Chiara Boldrini
IIT-CNR
Via G. Moruzzi 1
56124, Pisa, Italy
chiara.boldrini@iit.cnr.it

Marco Conti
IIT-CNR
Via G. Moruzzi 1
56124, Pisa, Italy
marco.conti@iit.cnr.it

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

ABSTRACT

Message delivery in opportunistic networks is substantially affected by the way nodes move. Given that messages are handed over from node to node upon encounter, the intermeeting time, i.e., the time between two consecutive contacts between the same pair of nodes, plays a fundamental role in the overall delay of messages. 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 study of the conditions for the expectation of message delay to converge when social-oblivious forwarding schemes are used. More specifically, we consider different classes of social-oblivious schemes, based on 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. Our main finding is that, as long as the convergence of the expected delay is concerned, allowing more than two hops does not provide any advantage. At the same time, we show that using a multi-copy scheme can, in some cases, improve the convergence of the expected delay.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication*; C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing protocols*

Keywords

opportunistic networks, forwarding protocol, expected delay convergence

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*, typically 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 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 [13]. Forwarding protocols may also differ in the number of relays that they exploit. Simpler strategies may be single hop or two hops strategies (e.g. Direct Transmission and Two Hop [7]), 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. As messages follow multi-hop paths across the nodes of the network, their delay is the result of the delay accumulated at each hop along the forwarding path. Therefore, the time (*intermeeting time*) between con-

MobiOpp'12, March 15–16, 2012, Zürich, Switzerland.

secutive encounters of a pair of nodes is the elementary component of the overall delay. Thus, knowing the distribution of intermeeting times, one could - in principle - model the distribution of the delay experienced by messages. Unfortunately, there is no agreement on the actual shape featured by pairwise intermeeting times in real networks. Of the many hypotheses that have been made [5] [6] [12][10], the most challenging from the forwarding standpoint is the one proposed by Chaintreau et al. [4]. Chaintreau et al. found intermeeting times extracted from real mobility traces to follow a Pareto[1] distribution. 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 might diverge. Clearly, having a finite expected delay is a key requirement for any communication protocol.

Due to the relevance of the problem, in this paper, considering heterogeneous mobility patterns, we derive the conditions on the Pareto exponent of pairwise intermeeting times under which the expectation of the delay under multi-copy and/or multi-hop social-oblivious protocols is finite. The starting point of our paper is the work by Chaintreau et al. [4], where such conditions have been studied for the the single-copy two-hop scheme and flooding (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). Recent works have shown, however, that real networks are intrinsically heterogeneous [5]. In this paper, we investigate whether heterogeneity in contact patterns helps the convergence of the expected delay of a general class of social-oblivious forwarding protocols and whether convergence conditions can be improved using multi-copy strategies and/or multi-hop paths.

We anticipate here that longer paths, i.e., with a number of allowed hops greater than two, do not help the convergence of social-oblivious strategies. The intuitive explanation is that two hops are enough for exploring the forwarding diversity available in the network. In fact, the relay to which the source hands over the message under the two-hop scheme can be any other node in the network, just as in the multi-hop case in which the number of allowed hop is greater than two. On the other hand, 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 finding the destination.

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 the main techniques that can be applied to social-oblivious forwarding, thus identifying a set of representative classes of social-oblivious schemes. For these classes, we derive in Section 5 the conditions for the expectation of their delay to be finite. Finally, Section 6 concludes the paper.

[1]In the following we use the terms "Pareto" and "power law" interchangeably.

2. RELATED WORK

As discussed above, 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 the following we only consider social-oblivious schemes, as they are the focus of this work. The simplest social-oblivious protocol is Direct Transmission [7], 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 [14] 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 unfeasible due to the huge amount of resources it consumes [13]. 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 [13], gossiping [8]). Another popular social-oblivious forwarding protocol is the Two Hop scheme [7], 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 throughput capacity in a homogeneous network [7].

To the best of our knowledge, there is no other contribution besides that of Chaintreau et al. [4] 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 realistic case of heterogeneous intermeeting times (as opposed to the homogeneous mobility considered in [4]) and we extend the set of social-oblivious policies considered. As a check of correctness, in Section 5 we apply our derivation to the homogeneous case and, under the same configuration used by Chaintreau et al. in [4], we obtain the same results.

3. NETWORK MODEL

Our model considers a network with N mobile nodes. We denote with $\mathcal{N}$ the set of all nodes in the network. 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 [3], 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).

Given that messages are handed over from node to node before reaching their destination, the way nodes move heavily affects the delay experienced by messages. As we assume that the transmission of a message can always be completed during a pair-wise 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. If t_f is the time at which a contact between node i and node j has* just *finished, the intermeeting time M_{ij} is given by:*

$$M_{ij} = \min_{t>t_f}\{t - t_f : ||X_i(t) - X_j(t)|| < r\} \qquad (1)$$

where $X_i(t)$ and $X_j(t)$ denote the position of i and j at time t, and r is the transmission range[2]*.*

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 $\frac{1}{E[M_{ij}]}$.

The message generation process and the mobility process are 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 (Residual Intermeeting Time). *Assuming that node i and node j are not in contact at time t_o, the residual intermeeting time $R_{ij}(t)$ between them is given by the time interval between t_o and the first time node i and node j come into each other's range again, i.e.:*

$$R_{ij} = \min_{t>t_o}\{t - t_o : ||X_i(t) - X_j(t)|| < r\}, \qquad (2)$$

where $X_i(t)$ and $X_j(t)$ denote the position of i and j at time t, and r is the transmission range.

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:

$$F_{M_{ij}}(t) = \left(\frac{t_{min_{ij}}}{t + t_{min_{ij}}}\right)^{\alpha_{ij}} \qquad (3)$$

in which we use the definition of the Pareto distribution which allows for values arbitrarily close to zero, usually denoted as American Pareto [11] [1] (as opposed to the European Pareto version). 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. Thus, the following remark holds.

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

[2]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.

N	number of nodes in the network
F_X	complementary cumulative distribution function (CCDF) of random variable X
$X(x)$	probability density function of random variable X
M_{ij}	intermeeting time for the i, j node pair
R_{ij}	residual intermeeting time for the i, j node pair
α_{ij}	exponent (*shape*) of the Pareto distribution that characterises M_{ij}; we assume $\alpha_{ij} > 1, \forall i, j$
t_{min}	scale of the Pareto distribution that characterises M_{ij}; $t_{min} > 0$
D_i^d	delay of a message generated by node i and addressed to node d
$\mathcal{N}$	set comprising all nodes of the network
$\mathcal{P}_i$	set comprising all nodes that can be encountered by node i
h_{max}	maximum number of hops allowed

Table 1: Notation

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 [1], 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 [1]. 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 $\alpha_{ij} > 2$.*

The notation used throughout the paper is summarised in Table 1. Similarly to the reference literature [4][9], for ease of computation in the following we restrict to the case of power law random variables having the same scale, i.e., $t_{min_{ij}} = t_{min}, \forall i, j \in \mathcal{N}$. In addition, for the sake of comparison with [4], 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 others. 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 therefore the distributions of intermeeting times.

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.

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. These multiple copies can be all created and handed over by the source node, or also intermediate relays could be allowed to take part into the multiplication process. Here we only focus on source generation. Other possible configurations (e.g., intermediate relays allowed to generate new copies, like in the Spray&Wait case [13]) 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., if each possible relay can be exploited just once, messages cannot perform more than $N-1$ hops). When the number of allowed hops is finite, the last relay can only deliver the message to the destination directly.

Third, the amount of knowledge that each agent in the forwarding process can rely on (or is willing to collect and store) is an additional element for classifying forwarding strategies. Focusing on the source node, there can be social-oblivious strategies in which the source node does not keep track at all of how the forwarding process progresses. In this case, considering the configuration in which the source node can generate up to m copies of the message, the m copies might end up being all distributed to the exact same relay, thus eliminating the potential benefits of multi-copy forwarding. A memoryful source, instead, is able to guarantee to use distinct relays. A similar problem holds for intermediate relays. Memoryless relays can forward the message to the same next hop more than once, because they are not at all aware of what happened in the past. On the other hand, memoryful relays possess this knowledge, and are able to refuse the custody of messages that they have already relayed. Please note that we assume that the source node can never be handed over messages that it has generated. This assumption simply takes into account the fact that the source identity is always enclosed into the message header, thus this does not require any additional knowledge beside what is already present in the system.

Table 2 summarizes the feasible combinations (the ones marked with the checkmarks) of the forwarding characteristics described above when social-oblivious schemes are considered. These combinations can be found in well known routing strategies. For example, the 1-hop 1-copy memoryless forwarding corresponds to the Direct Transmission strategy [7], in which the source node can only deliver the messages to the destination. The 2-hop 1-copy memoryless forwarding is equivalent to the Two Hop forwarding introduced in [7]. The 2-hop m-copy memoryful forwarding is equivalent to the multi-copy version of the Two Hop protocol studied in [4]. Please note that relays can be memoryful only when they have multiple forwarding choices. This is not the case when the number of hops is limited to either one (there is no relay in this case) or two (relays can only deliver the message to the destination).

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 would require combining the conditions derived in the paper, but the problem is not touched upon due to lack of space. 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 (i.e., the probability of an encounter with node i is strictly greater than zero). Recall, also, that we assume that $\alpha_{ij} > 1$ for all i, j node pairs, so that the residual inter-meeting times are defined (see Remark 2). 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. We refer the interested reader to [2] for the complete proof.

5.1 Single-copy schemes

In a previous work we have studied the single-copy case for both the 1-hop and 2-hop social oblivious forwarding protocols. For the readers' convenience, we hereafter recall these findings, whose proofs can be found in [2].

THEOREM 1 (SINGLE-COPY ONE-HOP SCHEME). *In a heterogeneous network where the intermeeting time M_{ij} between any generic i, j node pair follows a power law distribution with shape α_{ij}, when the Single-copy One Hop relaying protocol (also known as Direct Transmission 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 (SINGLE-COPY TWO-HOP SCHEME). *In a heterogeneous network where the intermeeting time M_{ij} between any generic i, j node pair follows a power law distribution with shape α_{ij}, when the single-copy two-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 $\sum_{j \in \mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$, *where $\mathcal{P}_s$ denotes the set of all nodes that can be encountered by node s;*

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

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. Recall that, as shown in Table 2, with the n-hop single-copy social-oblivious forwarding we must consider both the memoryless and the memoryful case for relays. Thus, in the memoryless case, relays hand over the message to the first encountered node,

	1 hop		2 hops		n-hop	
	1 copy	m copies	1 copy	m copies	1 copy	m copies
memoryless	✓	-	✓	✓	✓	✓
memoryful source	-	-	-	✓	-	✓
memoryful relays	-	-	-	-	✓	✓

Table 2: Summary of social-oblivious routing strategies

regardless of whether this node has already relayed the message or not. On the other hand, memoryful relays guarantee that the message is relayed at most once by each node.

THEOREM 3 (SINGLE-COPY n-HOP SCHEME). *In a heterogeneous network where the intermeeting time M_{ij} between any generic i, j node pair follows a power law distribution with shape α_{ij}, when the single-copy n-hop relaying protocol (both in the memoryless and memoryful case) 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. Here we only provide a sketch for the proof, whose complete version can be found in the associated technical report [2]. The proof is composed of three parts. We first study the delivery from the source node to the relay, then we concentrate on the delivery from relay to relay along the multi-hop path, and finally we study the delivery from the last relay to the destination node.

The source node s can either deliver the message directly to the destination or hand it over to an intermediate relay. Recall that we model message arrival time as a random point in time with respect to the evolution of the mobility process. Thus, the time before the source node releases 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. It can be showed 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 2, should be greater than 1 for the expectation to converge. This implies $\sum_{j \in \mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$, thus obtaining condition C1.

Once the source node has handed over the message, we know that the message will follow a n-hop path, with $n \leq \min\{N-1, h_{max}\}$ for the memoryful case and $n \leq h_{max}$ in the memoryless case. First, note that any node $i \in \mathcal{N} - \{s, d\}$ has a non negligible probability of being the k-th hop along the n-th hop path, with $k \in \{1, ..., n-1\}$. In fact, assume for a moment that the message can leave any node within a finite expected delay (conditions under which this assumption is true are provided below). Then, given that we assume $\alpha_{ij} > 1$ for all i, j node pairs, i.e., that nodes can meet with any other node, at each forwarding step every node has a non negligible probability of being selected.

Let us now derive the conditions for the expected time before the message leaves a node to be finite. Before considering intermediate relays, let us focus on the delivery from the last relay to the destination node. It is possible to prove that the delivery from the last relay to the destination shares the same convergence condition on its expectation as the residual intermeeting time between the relay and the destination. From Remark 2 we know that R_{jd} has finite expectation if $\alpha_{jd} > 2$. Given that all nodes have a non negligible probability of being the $(n-1)$-th hop, as we proved above, condition $\alpha_{jd} > 2$ must be satisfied for all nodes $j \in \mathcal{N} - \{s, d\}$. Under our assumption of nodes all potentially meeting with each other, this condition is equivalent to condition C2, as $\mathcal{N} - \{s, d\} = \mathcal{P}_s$.

In order to complete the proof we should also derive the conditions under which the delivery from intermediate relay to intermediate relay achieves finite expected delay. This derivation is quite involved, thus we left it entirely to the associated technical report. Note, however, that, given that conditions C1 and C2 are required for the first and last hop, the overall convergence conditions for the single-copy n-hop scheme can at most be equal to those of the single-copy 2-hop scheme, not better. Specifically, in [2] we derive that condition C2, that must apply to all nodes in $\mathcal{N} - \{s, d\}$, guarantees that the conditions derived for intermediate relays are automatically satisfied, both for the memoryless and the memoryful case. Thus, overall, conditions C1 and C2 guarantee the convergence of the expected delay under then n-hop single-copy social-oblivious scheme. □

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 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. In the following we present the results for the convergence of multi-copy/multi-hop schemes, while a detailed discussion of their advantages and disadvantages with respect to the single-copy schemes analyzed above will be provided in Section 5.3. Please also note that hereafter we only provide an intuitive sketch for the proofs, which can be found in a detailed version in the associated technical report [2].

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 met. Moreover, in the *memoryless* case, the source node does not keep a record of the relay nodes used so far, and thus two consecutive encounters with the same node will end up in the message being copied again to the same relay. In the *memoryful* case, a relay node cannot be used more than once. As we discuss below, these different capabilities have a great impact on the convergence of the expected delay.

THEOREM 4 (m-COPY MEMORYLESS TWO-HOP). *In a heterogeneous network where the intermeeting time M_{ij} between any generic i, j node pair follows a power law distribution with shape α_{ij}, when the memoryless multi-copy two-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 two parts. First, we discuss the convergence conditions for the m copies sent by the source node. The first copy can be studied analogously to what we did in Theorem 3, thus obtaining condition C1. For the following copies, we have to consider that relays can be reused, thus the time before the k-th copy leaves the source is given by the minimum of both residual intermeeting times (for nodes that have not been yet used as relays) and intermeeting times (for nodes already used as relays), because the mobility process regenerates upon encounter. We prove that the fact that the delivery of the k-th copy is constrained to start after all previous copies have been delivered does not affect convergence for the mobility scenario considered in the paper and, for the sake of clarity, here we neglect it (it is, however, considered in [2]). Thus, we have that, every time a new copy is handed over, a residual intermeeting time in the initial set $\{R_{sj}\}_{j \in \mathcal{P}_s}$ from which we take the minimum is substituted with the corresponding intermeeting time. Given that convergence conditions are looser for intermeeting times than for residual intermeeting times (see Remarks 1 and 2), it follows that condition C1, which is the necessary and sufficient convergence condition that applies to the case in which there are only residuals, is also a sufficient condition for the cases in which intermeeting times and residual intermeeting times are mixed.

The analysis of the second hop (delivery from relays to the destination) starts from the consideration that, given that the protocol is memoryless, the number of distinct relays actually carrying a copy of the message ranges from 1 to m. The worst case from the convergence standpoint is that in which all copies have been relayed to the same node. As this can happen with a non negligible probability, this worst case would impair convergence and thus we have to avoid it ensuring that condition C2 holds for all possible relays. □

In the following we derive the convergence conditions for the expected delay under the memoryful m-copy two-hop scheme. To this aim, in Lemma 1 we prove the existence of an operating point m^* for the memoryful m-copy two-hop scheme such that, when $m \leq m^*$, the expected time before all m copies are delivered to their m relays converges, while for $m > m^*$ copies exceeding m^* never achieve a convergent expected delay. Please recall that in this paper we assume $\alpha_{ij} > 1$ for all i, j node pairs.

Lemma 1. *In a heterogeneous network where the intermeeting times M_{ij} between any generic i, j node pair follow a power law distribution with shape α_{ij} and the memoryful m-copy two-hop forwarding protocol is in use, there exists a characteristic value m^* such that, when $m \leq m^*$, the expected time before all m copies are delivered to their m relays converges, while for $m > m^*$ copies exceeding m^* never achieve a convergent expected delay. The value of m^* can be obtained as follows:*

$$m^* = \begin{cases} 0 & \text{if } \sum_{j \in \mathcal{P}_s} \alpha_{sj} \leq N \\ \arg\max_m \{m + \sum_{i=m}^{N-1} \alpha_i^* > 1 + N\} & o.w. \end{cases}, \tag{4}$$

where α_i^* denotes the i-th largest α_{sj} with $j \in \mathcal{P}_s$.

Proof. We consider a memoryful source that is delivering the m initial copies of the message. As the source is memoryful, after the k-th copy is relayed, the next copy can be delivered to the subset of nodes that comprises only those that have not been already used as relay. We prove that the convergence conditions become stricter as the cardinality of the set from which we choose the relays decreases. This let 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 $N - m$. Following the same line of reasoning discussed in the proof of Theorem 3, we prove that all nodes have a non negligible probability of being chosen as relays at each step. Thus the sets of possible relays are given by all possible combinations of $N-m$ relays taken from the initial $N - 1$. The worst combination, as far as the convergence of the expected delay of the m-th copy is concerned, is that containing the $N - m$ nodes having the smallest Pareto exponent α_{sj}, with $j \in \mathcal{P}_s$. If we show that the m-th copy is handed over by the source node in a finite expected time, then the convergence for all previous copies and all cases different from the worst one will automatically follow. We derive that the m-th copy is relayed within a finite expected time if the sum of the $N - m$ smallest exponents α_{sj} with $j \in \mathcal{P}_s$ is greater than $1 + N - m$. In order to find the value m^* corresponding to the maximum number of copies that can be sent within a finite expected time, we simply compute the maximum m value such that the above condition is satisfied. □

Finally, in Theorem 5 we provide the convergence conditions for the overall expected delay under the memoryful m-copy two-hop scheme operating at $m \leq m^*$.

Theorem 5 (m-copy Memoryful Two-Hop). *In a heterogeneous network where the intermeeting times M_{ij} between any generic i, j node pair follow a power law distribution with shape α_{ij}, when the memoryful m-copy two-hop forwarding protocol operating at $m \leq m^*$ is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if $\sum_{j=N-m}^{N-1} \alpha'_j > 1+m$, where α'_j denotes the j-th largest α_{jd} with $j \in \mathcal{P}_s$ (thus $\sum_{j=N-m}^{N-1} \alpha'_j$ is the sum of the m smallest α_{jd} with $j \in \mathcal{P}_s$).*

Proof. The proof focuses on the second hop, as the delivery from source to relay is guaranteed to have finite expected delay by condition $m < m^*$. The second hop can be modelled as a parallel delivery from the m relays to the destination. For each of the relays, the residual intermeeting time with the destination starts from when the message has been received by the relay. In the worst case, the set of relays currently holding a copy of the message is composed by nodes having the lowest exponent value for the intermeeting time with the destination. Starting from these considerations and using Remark 2 and some Pareto properties derived in [2], we prove condition $\sum_{j=N-m}^{N-1} \alpha'_j > 1 + m$. □

As discussed before, Chaintreau et al. [4] 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 1 we verify that Theorem 5 confirms and extends the results in [4].

Corollary 1. *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 m-copy two-hop strategy ($m \leq m^*$) is used, the expected delay for messages generated by the source node s for the destination node d converges if and only if*

$$\alpha > \frac{1}{N-m} + 1. \quad (5)$$

In addition, m^* is given by $m^* = \left\lfloor N - \frac{1}{\alpha-1} \right\rfloor$.

Please note that the necessary and sufficient condition in Equation 5 extends the sufficient condition provided by Chaintreau et al. [4]. In fact, Chaintreau et al., under the assumption $N > 2m$ (which we have relaxed), derive that the expected delay of the m-copy two-hop scheme ($m \leq m^*$) converges in a homogeneous setting as long as $\alpha > 1 + \frac{1}{m}$. Exploiting assumption $N > 2m$, we have that $N - m > m$, thus $\frac{1}{N-m} < \frac{1}{m}$, and $1 + \frac{1}{N-m} < 1 + \frac{1}{m}$. Thus, when condition $\alpha > 1 + \frac{1}{m}$ is verified, also Equation 5 holds true.

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, the message travels along multi-hop social-oblivious paths until the destination is found. Based on the type of memory applied to the source node and to the relays, we consider the following versions of the n-hop m-copy protocol (corresponding to the last column of Table 2):

V1 the source node does not keep track of already used relays nor the intermediate relays do

V2 the source node selects m distinct nodes but the intermediate relays are not aware of already used relays

V3 the source node selects m distinct nodes and the intermediate relays can relay the message only once.

Theorem 6 describes the convergence conditions that apply in all these cases.

THEOREM 6 (m-COPY n-HOP). *In a heterogeneous network where the intermeeting time M_{ij} between any generic i, j node pair follows a power law distribution with shape α_{ij}, when either the V1, V2, or V3 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. Due to lack of space, here we only discuss the V3 case, as, being it memoryful, it is expected to perform better than memoryless social-oblivious forwarding. Proofs for V1 and V2 forwarding can be found in [2]. As we did before, here we only sketch the proof and we refer the reader to the associated technical report for the rigorous mathematical derivation.

With V3 forwarding, any time a copy of the message is handed over (either by the source or by an intermediate relay) a relay is removed from the set of possible relays. We identify a worst case (which we show to happen with non negligible probability) in which the first copy overtakes all other copies. This is the case in which the source node is not able to hand over the second copy because the first one has already used all possible relays. Thus, in the worst case V3 forwarding becomes a single-copy multi-hop forwarding, for which Theorem 3 holds. Please note that in V3 forwarding Lemma 1 does not hold due to the overtaking effect. □

5.3 Discussion

Table 3 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 3) 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 convergences 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 m^* copies of the message. In the worst case $m^* = 1$, and thus we find again conditions C1 and C2 that hold for the 1-copy 2-hop strategy. But if the source node can reach operating point $m^* > 1$, conditions on the delivery from the relays to the destination become less restrictive as condition C4 can tolerate exponents α_{jd} smaller than 2 (the exact tolerance depends on the actual value of m^*).

5.4 A case study

In this section we apply the convergence conditions derived above to a toy scenario that we generate in order to illustrate the main differences among the social-oblivious strategies as far as the convergence of their expected delay is concerned. More specifically, we focus on forwarding strategies with different convergence conditions, which, as shown in Table 3, are 1-hop 1-copy (Direct Transmision), 2-hop 1-copy, and 2hop m-copy (with memoryful source) schemes. We consider 10 nodes, and the following set of exponents:

$$\boldsymbol{\alpha} = \{2.1, 2, 1.9, 1.8, 1.7, 1.6, 1.5, 1.4, 1.3\},$$

whose components are denoted as $\alpha_1, ..., \alpha_{N-1}$. We assume that a generic node i meets all other nodes in a way such that $\alpha_{i,1} = \alpha_1, ..., \alpha_{i,i-1} = \alpha_{i-1}, \alpha_{i,i+1} = \alpha_i, ..., \alpha_{i,N} = \alpha_{N-1}$. We also set t_{min} to $1s$. Let us consider messages sent by source node 1 with destination node 10. According to the results of Section 5.3, in this case the expected delay for the Direct Transmission is not defined, because $\alpha_{1,10} = 1.3$, while it should be greater than 2 for convergence. Analogously, the convergence condition for the 1-copy 2-hop scheme is

	1 hop		2 hops		n-hop	
	1 copy	m copies	1 copy	m copies	1 copy	m copies
memoryless	$\alpha_{sd} > 2$	-	[C1,C2]	[C1,C2]	[C1,C2]	[C1,C2]
memoryful source	-	-	-	[C3,C4]	-	[C1,C2]
memoryful relays	-	-	-	-	[C1,C2]	[C1,C2]

Table 3: Summary of convergence conditions for social-oblivious routing strategies (C1 and C2 are defined in Theorem 2, $C3 = m \leq m^*$ and $C4 = \sum_{j=N-m}^{N-1} \alpha'_j > 1 + m$)

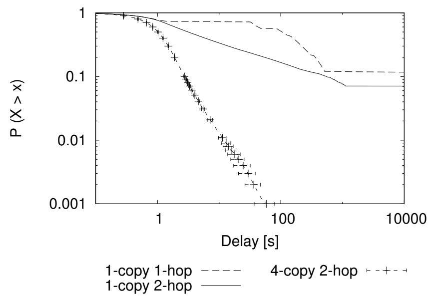

Figure 1: Comparison of delay CCDFs for the 1-hop 1-copy, 2-hop 1-copy, and 2-hop m-copy schemes

not satisfied. More specifically, condition C2 is not satisfied, because $\alpha_{i,d} < 2$ for all nodes i. The only scheme able to achieve a convergent expected delay is the m-copy 2-hop scheme. In fact, applying Lemma 1 we derive that the source can send up to $m^* = 4$ copies of the message for which the expectation of the latency is defined. If we assume to send all these four copies ($m = m^*$), condition C4 in Table 3 becomes $\sum_{j=6}^{9} \alpha_j > 5$. Given that the sum of the four smallest exponent in $\boldsymbol{\alpha}$ is equal to 5.8, condition C4 is satisfied.

In order to complement these results, we ran a set of simulations, using a custom simulator written in C++, in which node 1 sends messages to node 10 according to a Poisson process with mean 1 second. In order for the comparison to be fair, we run $20000s$ of simulated time and we considered only the messages generated in the first $10000s$ in our statistics. The $10000s$ packet lifetime has been chosen in order to be significantly greater than the expected delay ($\sim 2s$) from node 1 to node 10 when the 4-copy 2-hop scheme is used. Please recall that the 4-copy 2-hop protocol is the only social-oblivious scheme to achieve a finite expected delay in this scenario. When applicable, i.e., when the average value is finite, we also show the 99% confidence intervals. For the three forwarding strategies discussed above, we plot the empirical cumulative distribution function in Figure 1. As expected, in the case of 4-copy 1-hop scheme, the great majority of messages ($\sim 99.9\%$) is delivered within a short time ($100s$) from their generation. For both the 1-hop 1-copy and the 2-hop 1-copy schemes, instead, after 10000 seconds there is still a big fraction (around 10%) of messages to be delivered. These long delays, predicted by our model, are those that cause the expected delay to diverge.

6. CONCLUSIONS

Assuming heterogenous, Pareto distributed, intermeeting times, in this paper we have derived the conditions on the Pareto exponents such that the expected delay of a large family of forwarding protocols is finite. We have considered different classes of social-oblivious strategies based on the number of copies and the number of maximum relays that are allowed. Our main result is that convergence is not improved by an increased number of allowed hops. Specifically, there is no advantage, as far as the convergence of the expected delay is concerned, in using more than two hops. In addition, when the source node is able to directly deliver the message to the destination with a finite delay, any additional relay can only add more restrictive convergence conditions.

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. More specifically, a multi-copy two-hop strategy can prove effective when neither the source node nor intermediate relays are able to directly deliver the message to the destination within a finite expected time. The use of multiple copies, in fact, benefits from the parallel delivery of the message from different nodes, which may overcome the individual limitations in achieving a finite expected delay.

7. ACKNOWLEDGMENTS

This work was partially funded by the European Commission under the SCAMPI (FP7-FIRE 258414), RECOGNITION (FP7 FET-AWARENESS 257756), and EINS (FP7-FIRE 288021) projects.

8. REFERENCES

[1] C. Boldrini, M. Conti, and A. Passarella. From pareto inter-contact times to residuals. *IEEE Commun. Lett.*, 15(11):1256 – 1258, 2011.

[2] C. Boldrini, M. Conti, and A. Passarella. Less is more: long paths do not help the convergence of social-oblivious forwarding in opportunistic networks. Technical report, IIT-CNR, 2011, `http://cnd.iit.cnr.it/chiara/pub/mobiopp12_tr.pdf`.

[3] 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.*, 41(6):128–136, 2003.

[4] 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.*, pages 606–620, 2007.

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

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

[7] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. on Netw.*, 10(4):477–486, 2002.

[8] Z. Haas and T. Small. A new networking model for biological applications of ad hoc sensor networks. *IEEE/ACM Trans. on Netw.*, 14(1):27–40, 2006.

[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.*, 9(10):1377 –1390, 2010.

[10] I. Rhee, M. Shin, S. Hong, K. Lee, S. Kim, and S. Chong. On the levy-walk nature of human mobility. *IEEE/ACM Trans. on Netw.*, 19(3):630–643, 2011.

[11] M. Rytgaard. Estimation in the Pareto distribution. *ASTIN Bulletin*, 20(2):201 – 216, 1990.

[12] M. Seshadri, S. Machiraju, A. Sridharan, J. Bolot, C. Faloutsos, and J. Leskove. Mobile call graphs: beyond power-law and lognormal distributions. In *ACM SIGKDD'08*, pages 596–604. ACM, 2008.

[13] T. Spyropoulos, K. Psounis, and C. Raghavendra. Efficient routing in intermittently connected mobile networks: The multiple-copy case. *IEEE/ACM Trans. on Netw.*, 16(1):77–90, 2008.

[14] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical report, 2000.