

# Coordinating Replication Decisions in Multi-copy Routing for Opportunistic Networks

Nikolaos Papanikos and Evangelos Papapetrou

Department of Computer Science and Engineering, University of Ioannina, Greece

Email: {npapanik, epap}@cs.uoi.gr

**Abstract**—One of the dominant approaches in coping with the intermittent connectivity of opportunistic networks is packet replication. However, the need for a distributed operation forces the nodes carrying a message copy to make replication decisions without taking into account the replication state of other nodes. This strategy can lead to the creation of an excessive number of replicas thus exhausting the limited node resources such as energy and storage capacity. In this paper, we propose a simple yet efficient method which, without incurring any additional cost, allows nodes to share information about the replication process in order to avoid unnecessary replication. Our approach significantly reduces the routing cost without sacrificing delivery rate. Furthermore, it is generic in the sense that it can be implemented regardless of the utility metric used for making the replication decision. We validate the performance gains of our algorithm through analysis as well as extensive simulations.

**Index Terms**—delay tolerant networks, mobile opportunistic networks, multi-copy routing

## I. INTRODUCTION

Multi-copy strategies have been extensively used [1]–[8] in the context of routing in opportunistic and/or delay tolerant networks (DTNs) in order to tackle the problem of intermittent connectivity, i.e. the absence of end-to-end paths. The idea is simple; spreading more replicas increases the probability that a node carrying the packet will meet, i.e. move into the communication range of, the destination. Although this strategy achieves high performance in terms of delivery rate and delay, this comes at the cost of more transmissions and increased storage requirements. To illustrate this, let us consider the two extremes; on one hand stands the case that no replication occurs while on the other hand Epidemic routing [9] produces replicas in a greedy manner that exploits all contacts among nodes. Clearly, the first strategy suffers poor delivery rate since the message is delivered only if the source node meets the destination. On the contrary, Epidemic routing increases the delivery probability but is not suitable for a context of limited resources because it results in energy depletion and memory starvation at nodes. Controlling the level of replication allows for a tradeoff between delivery rate and cost (both energy and storage related). To this end, utility-based replication [6], [8] is probably the most appealing strategy due to its capacity to adjust to diverse network characteristics. The idea here is to introduce a utility metric that captures the fitness (or quality) of a node for delivering and/or forwarding the message and then create replicas by comparing the utility metrics of the nodes in contact. There exists a diverse range of metrics [1]–[3], [6], [7], [10]–[12] that are constructed from a node’s feature such

as the frequency or the regularity of its contacts, its importance in a social context, etc. Although the choice of the utility metric significantly impacts performance, it is common ground that, regardless of the metric used, utility-based replication frequently involves a high cost due to increased message replication. In order to reduce replication without significantly impacting the delivery rate, Delegation Forwarding (DF) [1] exploits the knowledge about past replication decisions by enabling each node to record the highest utility among its past contacts. This recorded value is the node’s perception of the highest utility in the network, therefore no replication is performed if the contacting node has a lower utility.

This work is motivated by the observation that message redundancy could further be reduced if each node makes replication decisions by taking into account not only its own perception of the highest utility in the network but also the highest utility as perceived by other nodes. In this way, nodes can coordinate their replication decisions in order to avoid unnecessary packet replications and thus achieve a more efficient replication process. We propose the Coordinated Delegation Forwarding (COORD) algorithm which incarnates the aforementioned functionality by allowing nodes to exchange their view’s regarding the replication process. We show that this can be done without any additional cost since the required information can replace the utility values that are usually exchanged during a contact. It is also important to note that COORD is generic in the sense that it can be implemented on top of any utility metric. We validate the performance gains of COORD through analysis as well as extensive simulations using traces from real networks. We show that it reduces the routing cost without sacrificing performance and that this is true for any utility metric.

The rest of the paper is organized as follows. In Section II, we review the work related to message replication. In Section III, we formulate the problem and discuss in detail our motivation. Then, in Section IV, we delineate the proposed protocol and provide an analysis of its routing cost proving its theoretical advantages. In Section V, we describe the simulation environment used in our experimental evaluation and discuss the simulation results while, in Section VI, we draw our final conclusions.

## II. BACKGROUND AND RELATED WORK

There exist two prominent approaches for replication in multi-copy routing schemes. In the first, the number of message replicas is bounded to a value  $L$  and the replicas are distributed

on a contact basis either blindly [8] or by ruling out nodes based on their fitness for delivering the message [6]. The downside of this approach is that defining  $L$  in order to achieve a certain tradeoff between delivery rate and cost is not straightforward. The second approach is more flexible in the sense that there is no requirement for predetermining the number of replicas to be created. Instead, each node’s ability to forward packets is determined by a utility metric and replicas are created based on the utilities of the encountering nodes. More specifically, suppose that node  $v$  meets node  $u$  and the corresponding utilities are  $U_v$  and  $U_u$ , respectively. There are two alternatives for replicating a message  $p$  from  $v$  to  $u$  [8]:

- *Absolute criterion:*  $U_u > U_{thresh}$ , for some threshold value  $U_{thresh}$
- *Relative criterion:*  $U_u > U_v + U_{thresh}$

Note that a utility may be either destination dependent, i.e. depend on  $p$ ’s destination, or independent. In the first case, a node  $v$  maintains one utility value for every possible destination  $d$  and the utility captures the ability of  $v$  to deliver packets to  $d$ . In the latter case, a destination independent utility captures the node’s ability to interact with other network nodes and consequently its fitness for acting as a forwarder regardless of the actual destination. Several utility metrics have been proposed, including the contact rate between nodes [3], [7], the time elapsed between node contacts [6], [10], the probability of node meetings [11], as well as metrics based on the social characteristics of nodes [2], [12].

Although the selection of the utility metric significantly impacts the performance of the algorithm, both the aforementioned criteria for performing replication tend to over-replicate messages, thus incurring an increased cost. To this end, Delegation Forwarding [1] introduces a replication criterion that takes into account the *history of a node’s observations*. More specifically, when node  $v$  meets node  $u$ , a packet  $p$  is replicated iff:

- *Delegation criterion:*  $U_u > \max_{k \in N_v} \{U_k\}$ , where  $N_v$  is the set of all nodes that  $v$  has contacted since the reception of  $p$ .

The rationale is that there is a little benefit in replicating a message to a node with a utility lower than the highest recorded utility so far. To further reduce the routing cost, Chen et al. [4] extend the delegation criterion by probabilistically pruning the set of possible message carriers. However, this approach can lead to performance degradation in terms of packet delivery. Furthermore, determining the optimal probability for ignoring a possible carrier of the message is a quite challenging task that strongly depends on the network characteristics. Gao et al. [5] present a work that also addresses the elimination of packet redundancy. However, this work considers only a subclass of destination independent utility metrics that are constructed as the sum of a node’s ability to reach every other network node.

This work builds on the premises of Delegation forwarding due to its efficiency but mostly due to its generic nature which allows it to be implemented with virtually any utility metric. We aim at further minimising the replication cost without however

sacrificing effectiveness by probabilistically suppressing replication. Instead, we take a deterministic approach and reduce the cost by taking advantage of the cooperation between nodes.

### III. PROBLEM FORMULATION AND MOTIVATION

To clearly demonstrate the motivation of this work, let us first examine the approach taken in [1], which is to model the replication process using the well-known problem of optimal stopping theory known as the hiring problem [13]. The problem concerns a small start-up company with the ambition to develop into a colossal and successful enterprise. To this end, the company interviews candidates in order to expand its work force and maximize the average employee quality. The only constraint is that the decision to hire or not a candidate must be instantaneous. According to this modeling, a node carrying a packet corresponds to the interviewer while the nodes that it contacts correspond to the possible employees. Furthermore, note that each node receiving a message copy may immediately also replicate this message. This, in the context of the hiring problem, is equivalent to the scenario where each chosen applicant immediately acts as a job interviewer and meets new applicants. In other words, the company’s hiring process contains multiple interviewers that evaluate candidates in parallel. One of the solutions to the hiring problem is known as the max strategy [13]. This strategy dictates that each candidate that has better quality from all current employees qualifies for a post in the company. Its main drawback is that it results in hiring a candidate more and more rarely as the time goes on [13]. However, in the aforementioned model, this drawback is eliminated by the parallel interviews which lead to a speed-up in the hiring rate [1]. This is why Delegation Forwarding (DF) implements a policy for producing replicas which resembles the max strategy, i.e. a node receives a copy if its utility is higher than the highest recorded utility so far.

We make the observation that DF’s policy on message replication deviates from the max strategy since it cannot maximize the quality of the nodes that carry a message copy. This is because, according to this policy, selecting a node for replication strictly depends on the node that already carries the message and its view of the maximum utility. Instead, the quality of message carriers could be improved if nodes exchange their view’s. This is equivalent to say that, in the context of the hiring problem, the interviewers are able to exchange information about the hired employees. COORD, without incurring any additional cost, exploits the recurring contacts between the message carriers in order to coordinate them to an up to date view of the highest utility seen in the network. The intuition is that this process will increase the quality of message carriers and as a result reaching the destination will require less replication.

### IV. COORDINATED DELEGATION FORWARDING

As mentioned previously, COORD works synergistically with any protocol that implements utility-based replication regardless of the utility that is used. Recall that in utility-based replication each node  $v$  is characterized by a utility  $U_v$  that

```

1: for every packet  $p \in Buf_v$  do
2:   if  $p \in Buf_u$  and  $c\tau_{v,t}^p < c\tau_{u,t}^p$  then
3:      $c\tau_{v,t}^p \leftarrow c\tau_{u,t}^p$ 
4:   else
5:     if  $c\tau_{v,t}^p < U_u$  then
6:       Forward  $p$  to node  $u$ 
7:        $c\tau_{v,t}^p \leftarrow U_u$ 
8:     end if
9:   end if
10: end for

```

Fig. 1. COORD's forwarding procedure when node  $v$  encounters  $u$  at time  $t$

portrays the node's fitness for delivering and/or forwarding a message. COORD introduces the concept of:

**Definition 1** (Coordinated Threshold:  $c\tau_{v,t}^p$ ). *The highest utility value among the nodes that carry a packet  $p$  which is known to node  $v$  at time  $t$ .*

Each node maintains a coordinated threshold for each packet that it carries. This threshold is initialized with the utility of the node, i.e.  $c\tau_{v,t}^p = U_v$ , when the message is received. Replication is performed by comparing the coordinated threshold of the node carrying the packet and the utility of the candidate node (Fig. 1, line 5). More specifically, in the event of a contact between nodes  $v$  and  $u$  at time  $t$ , a packet  $p$ , carried by  $v$  but not by  $u$ , is replicated only if  $c\tau_{v,t}^p < U_u$ , i.e. only if node  $u$ 's utility is higher than the highest utility known to  $v$ . Furthermore, the coordinated threshold is updated, i.e.  $c\tau_{v,t}^p \leftarrow U_u$ , since  $U_u$  is the new highest value known to  $v$ . The innovation of COORD is that a node  $v$ , when in contact with node  $u$ , is able to take advantage of the coordinated threshold of  $u$  for each packet that both nodes carry (Fig. 1, line 2), i.e.  $c\tau_{v,t}^p \leftarrow c\tau_{u,t}^p$  if  $c\tau_{u,t}^p > c\tau_{v,t}^p$ . This is done in order to allow  $v$  to increase its own coordinated threshold and therefore reduce future replication.

To illustrate the rationale behind COORD's approach, let us express more formally the update process of the coordinated threshold. Let  $K_v$  denote the set of node  $v$ 's contacts up to time  $t$ . Furthermore, let  $\langle u, T \rangle$  denote a contact of  $v$  with node  $u$  at time  $T$ . Then, the coordinated threshold of  $v$  for packet  $p$  can be expressed as:

$$c\tau_{v,t}^p = \max_{\forall \langle u, T \rangle \in K_v} \{c\tau_{u,T}^p, U_v, U_u\} \quad (1)$$

Note that  $U_v$  is the maximum value when the packet is first received by  $v$  while  $U_u$  is the maximum value when  $v$  replicates the message to  $u$  (in this case it is also  $c\tau_{u,T}^p = U_u$ ). Moreover, observe that the approach of DF is equivalent to the following update process:

$$\tau_{v,t}^p = \max_{\forall \langle u, T \rangle \in K_v} \{U_v, U_u\} \quad (2)$$

where  $\tau_{v,t}^p$  denotes the threshold in the case of DF. It is clear from (2) that in DF  $v$  uses only information regarding the utility of the nodes it meets. On the contrary, in COORD,  $v$  is able to exploit the utility of nodes that  $v$  have never been in contact with. This is possible through  $c\tau_{u,T}^p$  since it contains a "summary" of the utilities of  $u$ 's contacts which, in general, are different from  $v$ 's contacts. In other words, COORD exploits the recurring meetings between nodes in order to disseminate

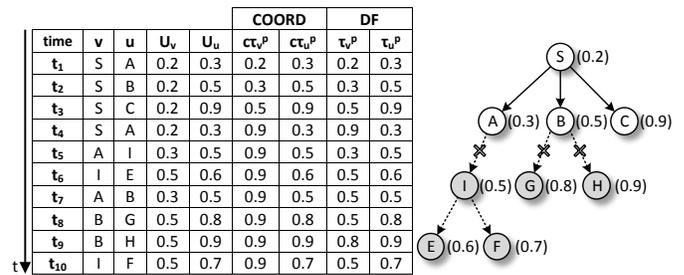


Fig. 2. Example of node meetings with the corresponding utility values and the replication tree

the highest utility seen across the network. To further illustrate the advantages of our approach, let us consider the example in Fig. 2. The figure depicts the replication tree constructed for a packet  $p$  which is originated at node S. The occurring contacts are illustrated in the corresponding table. At time  $t_3$  node C, which has a high utility value, receives  $p$ . DF uses this information to suppress future replication performed only by node S. On the other hand, COORD exploits the recurring contacts, at time instances  $t_4$  and  $t_7$ , to coordinate the views of nodes C, S, A and B. More specifically, A and B update their threshold values for packet  $p$  to 0.9 although they never meet C. As a result, the probability of performing future replication is also reduced for A and B. The outcome is a cost reduction of  $\sim 55\%$ . Finally, it is important to note that our approach comes at no additional cost, as no additional storage or transmissions of control packets are required compared to DF. The nodes in contact just need to exchange, for the packets that both carry, their coordinated thresholds instead of exchanging their utilities.

#### A. Cost Analysis

Similar to the plethora of multi-copy protocols we define the routing cost as the total number of packet replications that occur across the network. This definition captures both the number of transmissions, which is associated with node energy, as well as the packet load which is correlated to node storage requirements. Without loss of generality, in the following we focus on the case of routing a single packet  $p$ . Furthermore, we assume that for a node  $v$  both  $U_v$  and  $c\tau_{v,t}^p$  are normalized and take values in  $[0, 1]$ . Since DF is proved to have a lower cost than utility-based replication schemes [1], we focus on showing that COORD's cost is lower than that of DF. To this end, we first show that:

**Lemma 1.** *The coordinated threshold for a packet  $p$  at a node  $v$  is always equal or greater than the corresponding threshold of DF, i.e.  $c\tau_{v,t}^p \geq \tau_{v,t}^p, \forall t$ .*

*Proof.* We prove this Lemma by induction. Let  $T_0$  denote the time of the contact over which node  $v$  first receives  $p$ . Then, according to (1) and (2),  $c\tau_{v,T_0}^p = \tau_{v,T_0}^p = U_v^1$  since  $v$  receives a copy only if  $U_v$  is the highest value. Let  $T$  denote the time that the  $(k-1)$ -th contact occurs while  $T'$  denotes the time of the  $k$ -th contact and  $u$  denotes the contacting node with utility  $U_u$ .

<sup>1</sup>Hereafter, for simplicity, we omit the superscript  $p$  in the related notation.

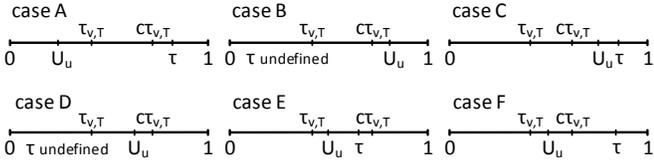


Fig. 3. Node  $v$  encounters node  $u$ : cases that result in an update of  $c\tau_{v,T'}$  and/or  $\tau_{v,T'}$

Furthermore, for simplicity, we use  $\tau$  to denote the threshold of  $u$  regardless of whether COORD or DF is used. Note that if  $u$  already has a copy of  $p$  then, by definition,  $\tau \geq U_u$  otherwise  $\tau$  is undefined. Since in both algorithms the threshold of a node is updated in a contact basis, it is sufficient to show that if  $c\tau_{v,T} \geq \tau_{v,T}$  then  $c\tau_{v,T'} \geq \tau_{v,T'}$ . Fig. 3 illustrates all the cases that, according to (1) and (2), result in updating  $c\tau_{v,T'}$  and/or  $\tau_{v,T'}$ :

**Case A:** Node  $u$  already has a copy of  $p$  and  $U_u \in [0, \tau_{v,T}]$  while  $\tau \geq c\tau_{v,T}$ . In this case the new thresholds are  $\tau_{v,T'} = \tau_{v,T}$  and  $c\tau_{v,T'} = \tau$ . In other words, the coordinated threshold is updated while the same does not happen for the DF case and as a result  $c\tau_{v,T'} \geq \tau_{v,T'}$ .

**Case B:** Node  $u$  does not carry a copy of  $p$  and  $U_u \geq c\tau_{v,T}$ . The new thresholds are  $c\tau_{v,T'} = \tau_{v,T'} = U_u$ .

**Case C:** Node  $u$  already has a copy of  $p$  and  $c\tau_{v,T} \leq U_u \leq \tau$ . After the update process  $\tau_{v,T'} = U_u$  and  $c\tau_{v,T'} = \tau$ , therefore  $c\tau_{v,T'} \geq \tau_{v,T'}$ .

**Case D:** Node  $u$  does not carry a copy of  $p$  and  $U_u \in [\tau_{v,T}, c\tau_{v,T}]$ . The update process will result in  $\tau_{v,T'} = U_u$  and  $c\tau_{v,T'} = c\tau_{v,T}$  and consequently  $c\tau_{v,T'} \geq \tau_{v,T'}$ .

**Case E:** Node  $u$  already has a copy of  $p$  and  $\tau_{v,T} \leq U_u \leq \tau \leq c\tau_{v,T}$ . After the update  $\tau_{v,T'} = U_u$  and  $c\tau_{v,T'} = c\tau_{v,T}$  and as a result  $c\tau_{v,T'} \geq \tau_{v,T'}$ .

**Case F:** Node  $u$  already has a copy of  $p$  and  $\tau_{v,T} \leq U_u \leq c\tau_{v,T} \leq \tau$ . The updated thresholds are  $\tau_{v,T'} = U_u$  and  $c\tau_{v,T'} = \tau$  which results in  $c\tau_{v,T'} \geq \tau_{v,T'}$ .

Note that in all of the aforementioned cases  $c\tau_{v,T'} \geq \tau_{v,T'}$ . Furthermore, in all other cases, the utility levels remain the same and consequently again  $c\tau_{v,T'} \geq \tau_{v,T'}$ .  $\square$

With the help of Lemma 1 we can now prove that:

**Theorem 1.** *The routing cost of COORD is lower or equal to the routing cost of Delegation forwarding.*

*Proof.* If  $N$  is the number of nodes in the network then the average cost for routing a single packet  $p$  is defined as:

$$Cost = [(N - 1) \cdot (1 - P_{nc})] + 1$$

where  $P_{nc}$  is the probability that a network node  $v$  will not receive a copy of  $p$ . This probability is equal to the probability that the utility of  $v$  is lower than the lower threshold among the nodes that  $v$  meets, i.e.:

$$P_{nc}^C = P\left[\min_{\forall \langle u, T \rangle \in K_v} \{c\tau_{u,T}^p\} \geq U_v\right]$$

in the case of COORD while in the case of DF:

$$P_{nc}^D = P\left[\min_{\forall \langle u, T \rangle \in K_v} \{\tau_{u,T}^p\} \geq U_v\right]$$

TABLE I  
PROPERTIES OF OPPORTUNISTIC TRACES

Trace Name	# Nodes	Duration (days)	Network Area
Infocom '05 [14]	41	3	conference
Sigcomm '09 [14]	76	3.7	conference
MIT Reality [15]	97	283	campus
Milano pmtr [14]	44	18.9	campus
Cambridge upmc [16]	52	11.4	city
Cabspotting [17]	511	2	city

Hence, it suffices to show that  $P_{nc}^C \geq P_{nc}^D$  or equivalently:

$$\min_{\forall \langle u, T \rangle \in K_v} \{c\tau_{u,T}^p\} \geq \min_{\forall \langle u, T \rangle \in K_v} \{\tau_{u,T}^p\}$$

which is directly derived from Lemma 1 since the lemma holds for all nodes.  $\square$

Note that the equality between the cost of COORD and DF holds only when a packet is replicated to nodes with increasingly higher utility even if the replication is performed by different nodes. It is clear that this is a rather unlikely case.

## V. PERFORMANCE EVALUATION

### A. Simulation Environment

In this section we evaluate the performance of COORD under various opportunistic environments. To this end, we developed a custom event-driven simulator that operates on a contact basis and is able to integrate real traces. We use three trace classes that differ in the network area in which the users move. The first class consists of two conference traces, the Infocom'05 and the Sigcomm'09 [14]. The second type is campus traces, where the participants are students and faculty members that move in a larger area than in the conference case. We use the well known MIT Reality dataset [15] and the Milano pmtr dataset [14], which utilizes a short beaconing scheme to get a more fine-grained view of contact records. The final trace category is the city-level traces in which the network area expands to the city limits. We have selected the upmc dataset [16] that was collected in the city of Cambridge, UK, and the Cabspotting dataset<sup>2</sup> [17] that consists of contacts between taxi cabs in San Francisco, USA. Table I summarizes the characteristics of the used traces.

As previously stated, the proposed algorithm is able to operate synergistically with a plethora of utility metrics. We assess the performance of COORD using various utility metrics both destination dependent and independent ones. More specifically, the following utilities are used:

**LTS [6]:** This is a destination dependent metric with values in  $[0, 1]$ . It is calculated as  $1/(1 + \text{LastTime})$ , where LastTime is the elapsed time since the last contact with the destination.

**DENC [7]:** This metric is calculated as the total number of contacts with the destination (destination dependent).

**ENC [3]:** This metric captures the total number of contacts with all network nodes (destination independent).

**BET [2]:** Betweenness Centrality is a destination independent metric that measures to what extent the node lies on the shortest paths from all nodes to all other. The distributed version of

<sup>2</sup>We extracted two days of the original dataset and processed them to generate node contact entries from the original GPS data.

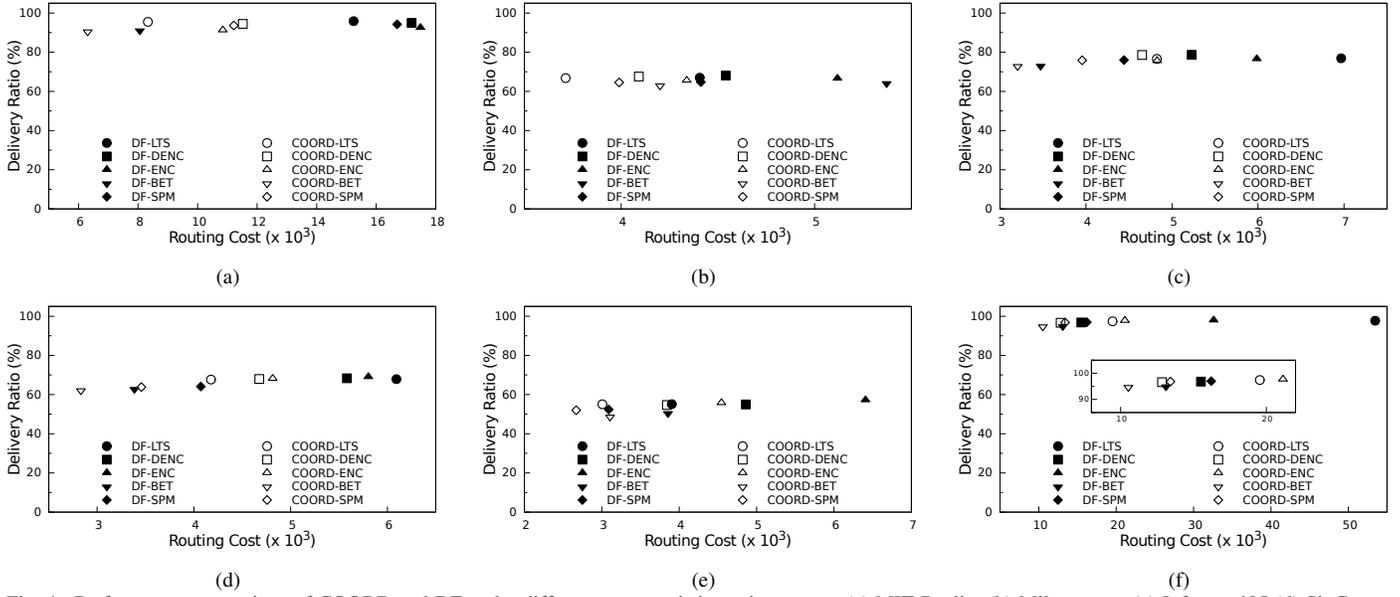


Fig. 4. Performance comparison of COORD and DF under different opportunistic environments: (a) MIT Reality (b) Milano pmtr (c) Infocom '05 (d) SigComm '09 (e) Cambridge upmc (f) Cabspotting

this utility, i.e. Ego Betweenness, is calculated using the local contact graph (i.e., ego network) of each node.

**SPM [12]:** Social Pressure Metric is a destination dependent metric that captures the friendship between network nodes. It is estimated locally by each node using the frequency, the longevity and the regularity of past node contacts.

The reported results are obtained as the average of 50 repetitions. Randomness is introduced in the selection of the source/destination pair for each packet as well as the time of packet generation. We randomly generate packets in the interval during which both the source and the destination are present in the network. Furthermore, we use a warm-up and a cool-down period, during which packets are not generated. The duration of each period is 20% of the total trace duration. In all cases, the confidence interval (with a 95% confidence level) of the reported results is less than 0.9%.

### B. Simulation Results

To evaluate the performance of our algorithm, we conducted two sets of simulations. In the first set, we examine the performance gains of COORD over DF under a diverse range of opportunistic environments. The traffic load is specified to

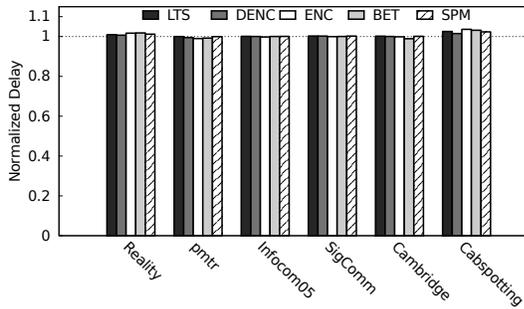


Fig. 5. Average packet delay of COORD, normalized to that of DF

1000 packets, each node storage capacity is unlimited and the number of replicas across the network is controlled using the Time-to-Live (TTL) deletion mechanism. The TTL value is set as the 20% of the total duration of each used trace. For the comparison, we use the following three metrics: the delivery ratio (i.e. the fraction of generated packets delivered to their destination), the routing cost (i.e. the total number of transmissions) and the average packet delay. Fig. 4 illustrates six plots that correspond to the six traces under evaluation. Each plot shows the delivery ratio versus the routing cost for both COORD and DF when various routing metrics are used. As expected, COORD clearly achieves a remarkable routing cost reduction compared to DF. More specifically, the gain ranges from 10% to 60% depending on the utility metric in use, as well as, the trace under consideration. Interestingly enough, the cost reduction comes at no or negligible impact on the delivery ratio. More specifically, when destination dependent metrics are used, the delivery ratio is virtually unaffected. In the case of destination independent metrics (i.e., BET and ENC), there is a minor impact on the delivery ratio, which is less than 1.77% in the worst case. To explain this, recall that COORD avoids replication to low utility nodes under the assumption that those nodes have a limited capability for delivering the message. However, when a destination independent metric is used this is not always the case, i.e. a low utility node may have increased delivery capability. However, this can hardly be considered as a shortcoming of COORD. It is actually a well-known disadvantage of destination independent metrics, which is also confirmed by the fact that those metrics consistently achieve the lowest delivery ratio in all traces (Fig. 4). Finally, Fig. 5 illustrates the average packet delay when COORD is used, normalized to that of DF. Clearly, COORD manages the routing cost reduction without significantly affecting the

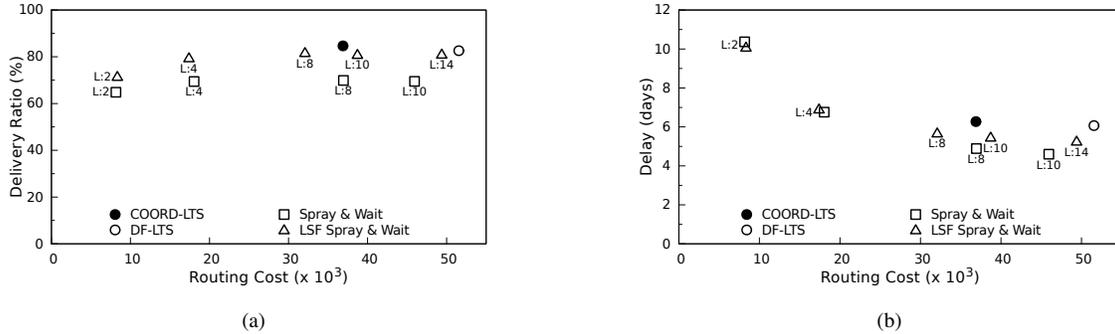


Fig. 6. Performance comparison of COORD and multi-copy routing schemes of bounded replicas: (a) delivery efficiency and (b) average packet delay

average delay for delivering packets.

In the second set of simulations, we assess the performance of our algorithm against multi-copy schemes that bound the message replicas to a predefined value  $L$ . In order to have a fair comparison, we choose COORD with the LTS utility, the LSF S&W [6] that also employs LTS and the S&W [8] which serves as a baseline algorithm for multi-copy schemes with a bounded number of copies. Furthermore, we limit the storage capacity of each node to 20 packets and generate 5000 packets in total. During a forwarding opportunity, packets are served using the FIFO policy and the oldest packet is discarded when the buffer is full. Fig. 6 presents the performance of the investigated protocols in the MIT Reality dataset. The results highlight the main drawback of both Spray-based algorithms which is the susceptibility of performance to the choice of  $L$ . Large values lead to good performance in terms of delivery and delay but at the expense of an excessive routing cost. Small values have the opposite effect and only a narrow range of  $L$ 's values provide a good tradeoff. Therefore, practically implementing spray-based protocols greatly depends on the non-trivial task of developing a highly accurate and self-configuring method for determining the suitable value for  $L$ . On the other hand, the proposed algorithm adapts to the network conditions without the need for predetermining the number of replicas. More specifically, COORD manages a gain in delivery ratio of  $\sim 5\%$  and  $\sim 15\%$  compared to the best performance of LSF S&W and S&W ( $L=8$  for both schemes). The slightly increased delay is mainly attributed to the bias introduced due to the increased delivery ratio, i.e. delivering packets to distant destinations is more difficult, therefore increasing the delivery ratio incurs longer delays.

## VI. CONCLUSION

We proposed a generic algorithm that can be implemented in synergy with virtually any utility-based replication scheme in opportunistic networks. The algorithm allows nodes to coordinate their views regarding the replication process and in this way achieves to significantly reduce the number of replicas without sacrificing the delivery rate and the end-to-end delay. The analysis as well as the extensive simulations in a diverse range of opportunistic environments have proved the performance gains over delegation forwarding, which is one of the state-of-the-art algorithms in this category. Finally,

we demonstrated the ability of our algorithm to adapt to the network characteristics and operate without the need of predetermining an upper bound on the number of replicas.

## REFERENCES

- [1] V. Erramilli, M. Crovella, A. Chaintreau, and C. Diot, "Delegation Forwarding," in *Proc. ACM Int. Conf. Mobile Ad Hoc Netw. and Comp. (MobiHoc)*, 2008, pp. 251–260.
- [2] E. M. Daly and M. Haahr, "Social network analysis for information flow in disconnected delay-tolerant MANETs," *IEEE Trans. Mobile Comput.*, vol. 8, no. 5, pp. 606–621, 2009.
- [3] S. C. Nelson, M. Bakht, and R. Kravets, "Encounter-based routing in DTNs," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, 2009, pp. 846–854.
- [4] X. Chen, J. Shen, T. Groves, and J. Wu, "Probability delegation forwarding in delay tolerant networks," in *Proc. IEEE Int. Conf. Comput. Commun. and Netw. (ICCCN)*, 2009, pp. 1–6.
- [5] W. Gao, Q. Li, and G. Cao, "Forwarding Redundancy in Opportunistic Mobile Networks: Investigation and Elimination," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, 2014, pp. 2301–2309.
- [6] T. Spyropoulos, T. Turletti, and K. Obraczka, "Routing in delay-tolerant networks comprising heterogeneous node populations," *IEEE Trans. Mobile Comput.*, vol. 8, no. 8, pp. 1132–1147, 2009.
- [7] V. Erramilli, A. Chaintreau, M. Crovella, and C. Diot, "Diversity of forwarding paths in pocket switched networks," in *Proc. ACM SIGCOMM Internet Meas. Conf. (IMC)*, 2007, pp. 161–174.
- [8] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: the multiple-copy case," *IEEE/ACM Trans. Netw.*, vol. 16, no. 1, pp. 77–90, 2008.
- [9] A. Vahdat, D. Becker *et al.*, "Epidemic routing for partially connected ad hoc networks," Duke University, Tech. Rep., CS-200006, 2000.
- [10] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli, "Age matters: efficient route discovery in mobile ad hoc networks using encounter ages," in *Proc. ACM Mobile Ad Hoc Netw. and Comput. (MobiHoc)*, 2003, pp. 257–266.
- [11] A. Lindgren, A. Doria, and O. Schelén, "Probabilistic routing in intermittently connected networks," *ACM SIGMOBILE Mobile Comput. and Commun. Rev.*, vol. 7, no. 3, pp. 19–20, 2003.
- [12] E. Bulut and B. K. Szymanski, "Exploiting friendship relations for efficient routing in mobile social networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 12, pp. 2254–2265, 2012.
- [13] A. Z. Broder, A. Kirsch, R. Kumar, M. Mitzenmacher, E. Upfal, and S. Vassilvitskii, "The Hiring Problem and Lake Wobegon Strategies," in *Proc. ACM-SIAM Symp. Discrete Algorithms*, 2008, pp. 1184–1193.
- [14] "CRAWDAD data sets: Infocom 05 (v. 2006-11-14), SigComm 09 (v. 2012-07-15), unimi-pmtr (v. 2010-09-10)," Downloaded from <http://crawdad.org>, Mar. 2014.
- [15] N. Eagle and A. Pentland, "Reality mining: sensing complex social systems," *Pers. and Ubiqu. Comput.*, vol. 10, no. 4, pp. 255–268, 2006.
- [16] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft, "Opportunistic content distribution in an urban setting," in *Proc. ACM SIGCOMM Workshop Challenged Netw. (CHANTS)*, 2006, pp. 205–212.
- [17] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "A Parsimonious Model of Mobile Partitioned Networks with Clustering," in *Proc. Int. Conf. Commun. Syst. and Netw.*, 2009, pp. 1–10.