Utility-based Forwarding: a Comparison in Different Mobility Scenarios

Elena Pagani
Computer Science Dept., Università degli Studi di Milano, Italy
CNR-IIT, National Research Council, Italy
pagani@dico.unimi.it

Gian Paolo Rossi
Computer Science Department
Università degli Studi di Milano, Milano, Italy
rossi@dico.unimi.it

ABSTRACT

Several proposals are available in the literature that deal with the problem of message forwarding in Opportunistic Networks (ONs). These proposals attempt to derive the path from source to destination that minimizes delivery latency and traveled hops, and maximizes the probability of successful delivery, while saving the overall system resources through a limitation of the number of message copies. Utility-based forwarding achieves these goals through the use of functions that discriminate among nodes in terms of their *utility* to reach a destination. Although the approach is very promising, so far, there is no understanding about the tight relationship between utility functions and the mobility scenario in which they operate and, as a consequence, we are unable to design efficient solutions for practical ONs.

In this work, we focus on this point by analysing five well known utility functions in five different scenarios. We establish relationships between the mechanisms adopted by the utility functions to discriminate among candidate relays, and the characteristics of the environment in terms of people mobility and the structure of their communities. The results can be useful to select an appropriate forwarding mechanism when deploying an experimental Opportunistic Network, and to design a novel utility function able to adapt to variable mobility patterns.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Network Architecture and Design

Keywords

Opportunistic networks, Utility-based forwarding

1. INTRODUCTION

The recent advances in short range radio technologies are driving the search for new wireless networking platforms that complement the 3G/4G network infrastructure with the aim of offloading cellular networks and providing a flexible alternative to deliver emerging mobile computing services such as, those location-sensitive (i.e targeted advertising, recommending systems) and contact-sensitive (i.e. mobile social networking, content sharing, urban sensing) [23, 12]. In this emerging heterogeneous networking scenario, Opportunistic Networks (ONs) [20] have great potential as a viable solution to enable the communications between the content source (either a mobile user in the neighborhood or a road-side AP) and the target mobile user(s). This communication is obtained by deploying a multi-hop path on top of contacts amongst mobile devices. The practical feasibility and the efficiency of such a path mainly depend on the function of forwarding.

The research has longly studied forwarding in ONs and proposed a variety of solutions in the literature [1, 2, 7, 8, 9, 10, 14, 16, 18]. For evident motivations, the most viable and practical solutions are those following the single-copy approach or a controlled multi-copy. The challenge is to approximate the performance and efficiency of the unviable, but optimal, forwarding achieved by using an oracle knowing all future contacts. Several interesting and practical solutions are available that attempt to approximate the oracle decisions by assigning utility values to nodes on the base of the available set of past and current contacts. If the values are properly assigned the forwarding is more likely able to discriminate, among the set of contacts, the relays that belong to the optimal path towards destination. The bet is that the past behavior of a node will be maintained in the future. In order to effectively discriminate among relays of different quality, contact dynamics should show two relevant behaviors. First, the contact pattern *within* every pair of nodes needs to be somehow regular. Irregular encounters make impossible to forecast the node likelihood of encountering the destination. Secondly, contact dynamics *between* different pairs of nodes need to be somehow heterogeneous, otherwise any node is an equally good candidate for data relaying.

Social and human sciences tell us that the human attitude of grouping in social communities, of sharing common locations or of commuting from one location to another, actually generates the required regularities and heterogeneities [21, 19, 22, 6]. This makes the approach promising but concentrates all challenges on the choice of the utility function. A good utility function captures the proper behavior of a given mobility pattern and is thus able to discriminate between good relay nodes (with more chances to encounter the destination in the future) and the others. A bad utility func-

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

tion is unable to do this and considers all nodes as equally useful.

For the above arguments, it is easy to argue that utility-based forwarding is highly influenced by the nature of underlying mobility patterns. Despite that, nobody has analyzed utility functions in different mobility scenarios. As a consequence, today we are unable to select the utility function that better fits with the nature of a given practical, deployable networking scenario or to design a good utility function for a given mobility setting. This paper focuses on these still open points by analyzing 5 well known utility functions in 5 different mobility scenarios. The results of the paper should be useful when deploying an experimental ON or to design a novel utility able to adapt to changing mobility patterns.

This is not the only paper approaching this argument; in [9], a subset of the utility functions we consider here have been analyzed; however, the analysis covers a very short time interval (3 hours) and uses real traces reproducing densely populated environments only.

2. MOBILITY SCENARIOS

ONs are supposed to be active in the mobile periphery of the Internet, the belt centered around the locations where users live, work or socialize. We can assume that each location covers a limited area (let us say at most 1000 × 1000 m.), involves at most a few hundred people and is characterized by mobility patterns that reflect the type of sociality people have there. In fact, the social destination of a place influences the way communities are formed, how people interact and move from group to group etc., thus conditioning the behavior of regularities and heterogeneities. With the intention of describing relevant human mobility and social attitudes we devised 5 scenarios that are described in the following.

The first scenario describes human attitudes in workplaces. Individuals in such a scenario are highly sedentary, but, due to the limited space, sooner or later also individuals belonging to separate communities (let us call them unfamiliar) happen to encounter one another. The mobility setting is characterized by contacts among familiar people lasting for long time and by unfrequent contacts between unfamiliar nodes. By slightly varying the previous scenario we can represents two new conditions that reproduce the human attitudes in spaces with higher mobility. In the former, very closed communities are considered. Sporadically, an individual belonging to a community may temporarily visit a different community. By contrast, a few individuals (namely, the *travelers*) continuously go back and forth thus having frequent contacts with people belonging to different communities. In the second, the community boundaries are softened and people belonging to different communities have more chances to encounter one another. In both cases, the contact duration is shorter than before – due to higher mobility – and the amount of global contact opportunities is higher. Two extreme conditions are also considered: a random and a deterministic mobility scenario. The latter yields perfectly predictable contacts; similar conditions can be found, for instance, in mobility patterns of a public transportation system.

2.1 Implementation of the Scenarios

We reproduce the five scenarios above by means of, respectively, a real trace, a synthetic model (HCMM [3]), and two benchmark models (Random Waypoint and a perfectly deterministic encounter pattern). We adopted HCMM because it can model human mobility patterns according to the influence of both popular locations and social relations. This is in contrast with, for instance, [13, 5, 17] that only consider one of the two aspects. Moreover, the setting of the HCMM parameters is very simple, easier than, for instance, in [15].

PMTR traces.

The real traces are obtained through a campus experiment and are generated by 44 people equipped with wireless devices, named PMTRs, with 10 m. radio range [11].[1] In order to compare the trace from a real dataset with those obtained from synthetic models, we eliminated nights and weekends from the experimental dataset, thus producing a dataset covering 13 working days, from 8:00 AM to 8:00 PM (156 h.). These samples are more than the 80% of the samples obtained from the whole experiment.

HCMM model.

In the two scenarios produced with HCMM, 44 nodes move in a 1000 × 1000 m. area with speed in [0.5, 1.5] m/s for 156 hours; the transmission range is 10 m. As an initial interaction matrix, we used weights derived from the number of contacts between pairs of nodes in the PMTR trace. The highest number of contacts was assigned weight 0.9. The weight associated to half of the average number of contacts has been adopted as a threshold to derive the connection matrix. No reconfiguration is performed and the remaining probability is set to 0.8. In the scenario named *HCMM_det5*, the next cell is chosen deterministically and the rewiring probability is 0.1; we adopted 5 travelers. In the *HCMM_pro1* scenario, the probabilistic criterion is adopted, with rewiring probability of 0.3 and 1 traveler.

Random Waypoint.

We produced a RWP scenario with BonnMotion v1.5 [4], with 44 nodes moving for 156 hours over a 500 × 500 m. area, with speed varying in the interval [0.5, 1.5] m/s. We adopted a pause interval of 3600 s. The trace produced by BonnMotion is post-processed in order to derive the contacts between nodes, assuming that the communication range is 10 m. and with sampling granularity of 1 s.

Deterministic model.

A deterministic (DET) scenario has been produced in the same environment as RWP as follows: communities of a 10% of nodes each have been set. Then, for each pair of nodes i and j, contacts are generated with parameters drawn from the mean values obtained with PMTRs (Table 1). The first contact occurs at a random time within the first hour. All contacts occur with a fixed inter-contact time ICT c_{ij} from one another and have a fixed contact length l_{ij}. In our experiments, c_{ij} has been extracted with uniform distribution within the range [1000, 1500] s. for nodes in the same community, and [5000, 6000] s. otherwise, while l_{ij} always varies

[1]The data set can be downloaded from the CRAWDAD archive (`http://www.crawdad.org/unimi/pmtr`).

Table 1: Comparison between synthetic scenarios and PMTR traces

scenario	# contacts	length (s.)	ICT (s.)	E(σ(ICT$_{ij}$))	σ(E(ICT$_{ij}$))	#hops $\mathcal{O}$	latency $\mathcal{O}$	%dest $\mathcal{O}$
DET	146.83	452.12	3369	0	1250	3.41	323 s.	100
PMTR	15.59	521.98	3381	3709	4198	3.22	32207 s.	87.8
HCMM_det5	292.28	11.88	1599	14738	21922	4.08	4367 s.	87.3
HCMM_pro1	58.75	11.94	8814	17243	20527	3.69	1273 s.	100
RWP_3600	22.05	123.03	24149	23674	6127	3.54	2130 s.	99.7

in the range [300, 600] s. This will allow to differentiate nodes in terms of their utility to reach a certain destination.

In both DET and RWP, all pairs of nodes eventually have a contact; this is not true in the other settings. However, in DET, nodes in the same group are more used to encounter one another, i.e. they encounter with a higher rate than with other nodes. In HCMM_det5, communities are quite closed, and their members mainly communicate through travelers. In PMTR, communities are not that closed, but people are sedentary and encounters occur seldom and are quite long. HCMM_pro1 still produces communities – unlikely RWP – but they are often mixed. As far as HCMM traces are concerned, we verified that the ICTs they produce fit with a Pareto's distribution. In Table 1, we show a comparison among the characteristics of the PMTR traces and those obtained by synthetic mobility models. All indices are averaged over all pairs of nodes and concern the whole time window. For PMTR, the ICT is evaluated only for consecutive contacts occurring within the same day. In columns 5 and 6 of Table 1, we report respectively: (*i*) the average on all pairs of nodes of the standard deviation of the ICT between node i and node j; (*ii*) the standard deviation among all pairs of nodes of the average ICT between node i and node j. The former gives an indication about the regularity of contacts *within* every pair. The latter gives an indication of the heterogeneity *between* different pairs of nodes. In Table 1, the performance achieved with an oracle-based optimal forwarding $\mathcal{O}$ is also reported, as the number of hops, latency, and percentage of reached destinations from each node to every other.

3. UTILITY-BASED FORWARDING

We study utility functions that capture various aspects of human interactions by using different mechanisms. In this work, we do not consider approaches involving community detection (e.g. [14]). Indeed, community detection is computationally expensive and still difficult to implement in a distributed manner. We do not assume that knowledge is a-priori available about movements, as in [18, 1]. We rather focus on mechanisms allowing to gain knowledge on contact dynamics in order to infer future encounters, e.g. through the analysis of the encounter history, as in [16, 9, 8], or the estimation of the centrality of nodes, as in [7]. This knowledge can be used for unicast communications. In particular, the following functions are considered:

- **Greedy (G) online:** the utility of a relay increases with the number of times it has encountered the message destination so far [9].
- **Greedy-total (GT) online:** the utility of a relay increases with the number of encounters (with any node) it has observed so far [9].

Table 2: Comparison between utility functions

	G	**GT**	**F**	**P**	**SB**
coloc. vs. social	coloc	coloc	coloc	coloc	social
dest. dependent	Yes	No	Yes	Yes	mix
aging	No	No	Yes	Yes	No
transitivity	No	No	No	Yes	No

- **Fresh (F):** a node $n1$ has greater utility than $n2$ for a destination D if the last encounter of $n1$ with D occurred more recently than that of $n2$ [8]. This approach has no memory of the past.
- **Prophet (P):** in [16], three mechanisms are used to maintain utilities. When two nodes $n1$ and $n2$ encounter, each one increments its delivery probability to the other as follows: $P(n1, n2) \leftarrow P(n1, n2) + (1 - P(n1, n2)) \cdot P_{init}$. A transitivity property of encounters is considered, such that, if $n2$ encountered $n3$, then the delivery probability of $n1$ to $n3$ is $P(n1, n3) \leftarrow P(n1, n3) + (1 - P(n1, n3)) \cdot P(n1, n2) \cdot P(n2, n3) \cdot \beta$. To take into account changing contact dynamics, an aging function is applied before each utility exchange such that $P(n1, n2) \leftarrow P(n1, n2) \cdot \gamma^t$, where t is the time elapsed since the last update.
- **SimBet (SB):** the utility of a node $n1$ for a destination D depends on both $n1$'s betweenness (i.e., whether it belongs to the shortest path between any two other nodes), and its similarity with D (i.e., whether there are nodes encountered by both $n1$ and D). We estimate both indexes as in [7].

In Table 2, we report a comparison among the characteristics of the considered approaches in terms of: whether colocation or social aspects are considered, whether the approach is destination dependent or not, whether it includes aging and transitivity mechanisms.

3.1 Measuring the Utility

In order to measure the goodness of a generic utility function $\mathcal{U}$, we consider its distance from optimal routing, able to follow the shortest path between source and destination.

To capture the dynamics of the utility values, we proceed as follows: every M minutes we freeze the utilities held at that time by nodes. According to those values we compute the ($\mathcal{U}$) path followed from each source to each destination, and we compute the distance between such path and the optimal one. Two indices Δ_e and Δ_l are considered to evaluate the distance and obtained as follows:

$$\Delta_e = \frac{\#hops\ of\ \mathcal{U}\ path\ -\ \#hops\ of\ \mathcal{O}\ path}{\#hops\ of\ \mathcal{O}\ path}$$

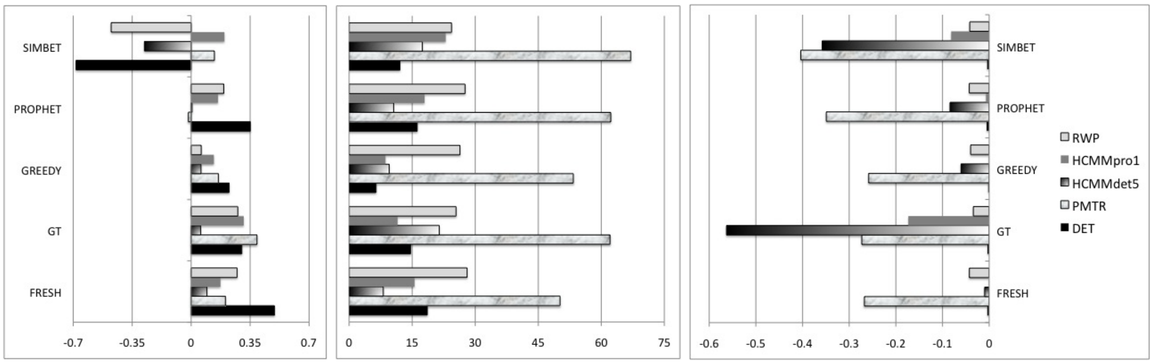

(a) (b) (c)

Figure 1: Comparison among approaches in terms of (a) Δ_e, (b) Δ_l, and (c) Δ_d.

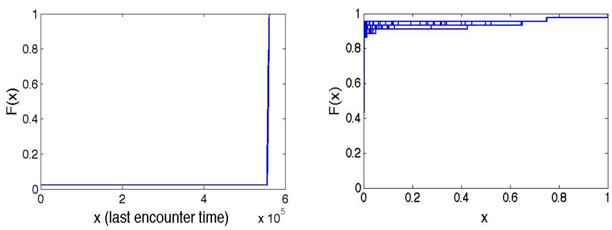

(a) (b)

Figure 2: DET: *ecdf* of the utility values (one plot for each node), for (a) Fresh, and (b) Prophet.

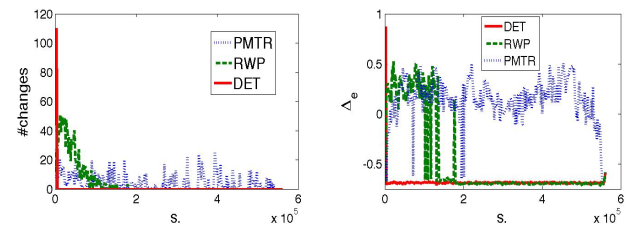

(a) (b)

Figure 3: SimBet: (a) utility changes, and (b) Δ_e vs. time.

$$\Delta_l = \frac{latency\ of\ \mathcal{U}\ path\ -\ latency\ of\ \mathcal{O}\ path}{latency\ of\ \mathcal{O}\ path}.$$

We use the path length in hops as an indirect measure of the overall system energy consumption (the higher the hop count, the more the devices involved to forward third party traffic over a radio channel). Δ_l measures how worse is the service offered to the users. Δ_e and Δ_l are measured just for the destinations reached by both $\mathcal{O}$ and $\mathcal{U}$. An index Δ_d is computed similarly, accounting for the fraction of destinations that $\mathcal{U}$ reaches with respect to $\mathcal{O}$. Clearly, if $\mathcal{U}$ follows paths with the same characteristics as $\mathcal{O}$, its distances will be 0.

4. SIMULATIONS

We used simulations to observe performance and behavior of the described utility functions in the different mobility scenarios. The following general settings have been adopted. The sampling rate is $M = 30$ minutes. This is also the length of the warm up period before generating the first message (this lets the utilities to initialize). The sampling is stopped after (7×12) hours leaving the last (6×12) hours for message delivery. Results are averaged over all source-destination pairs. The same single-copy forwarding policy is used for all approaches: if $n1$ has a message m for D and it encounters $n2$ with greater utility for D than its own, then $n1$ forwards m to $n2$ and removes m from its buffer.

Some more specific settings follow. Prophet simulations use the following parameters [16]: $P_{init} = 0.75$, $\beta = 0.25$, $\gamma = 0.98$. In SimBet, the parameter α, weighing between similarity and betweenness, is set to $\alpha = 0.5$ [7].

4.1 Results

In this Section, we focus on the simulation results with the aim of understanding the impact of different mobility settings on the performance of utility functions, and of capturing what mechanisms – among those the functions adopt – are successful to accurately discriminate among the utility of candidate relays. For a given mobility model, a good utility function is able to assign high values to nodes with more chances of encountering the destination in the future, while it assigns lower values to the other nodes. By contrast, a utility function that happens to flatten the assigned values cannot behave properly. Fig.1 summarizes the performance of the different approaches with the considered mobility traces. We now analyze in detail the relevant situations for each mobility model.

Let us consider the DET model first. The challenge in DET is identifying the groups of nodes that are used to encounter one another with the higher frequency, and assigning them a high utility. Here most of the approaches do not work well for slightly different motivations. All of them can hardly discriminate between nodes belonging to the community of the destination (that more likely meet the destination) and the other nodes (as shown by the *ecdf* of the utility values for the various destinations (Fig.2)). In Fresh, a node outside the community of the destination d might be chosen as forwarder because it encountered d very recently although it has few opportunities of being co-located with d. The performance of Prophet is negatively affected by the aging mechanism that flattens all the utilities to small values. As a consequence, the utility of nodes in the same community (*familiars*) is slightly higher than those of outside nodes (*un-*

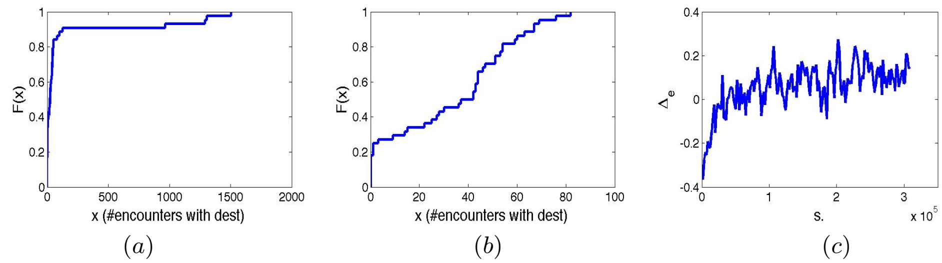

Figure 5: Greedy in HCMM_det5: ***ecdf*****(utilities) for (*a*) a sedentary node, and (*b*) a traveler. (*c*) Δ_e vs. time.**

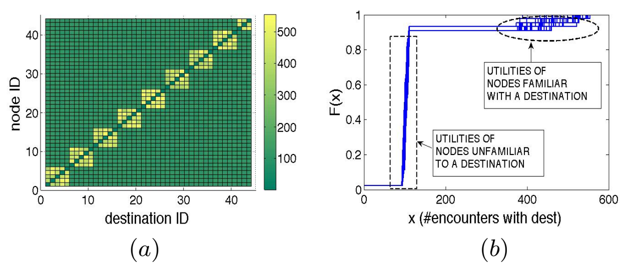

Figure 4: Greedy in DET: (*a*) utilities, and (*b*) ***ecdf*** **of the utility values (one plot for each node).**

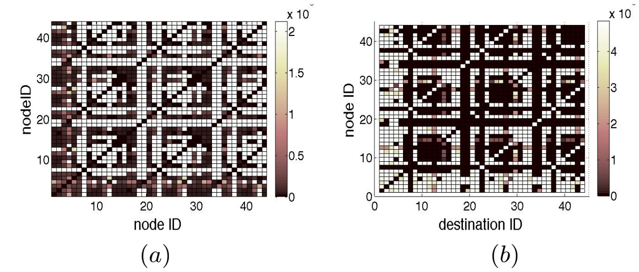

Figure 6: (*a*) ICTs with HCMM_det5, and (*b*) utilities learnt by Fresh.

familiars), but both are extremely low (0.107 and 0.018 on average respectively). Under these conditions, at each forwarding a very small progress towards the destination may be obtained, thus negatively affecting both latency and number of hops. A similar behavior is shown by Greedy Total: as the pattern of encounters is really similar for all nodes, utilities for the nodes are aligned and no node clearly distinguishes as preferred forwarder. An almost random choice does not guarantee that the destination is approached.

SimBet is more qualitative than quantitative: for each pair of nodes it records whether they have encountered or not. In DET (but similar arguments apply to RWP), when all nodes have encountered one another at least once, then no node is between any source-destination pair, and every node is equally similar to each other (both have encountered all the other nodes). At this point, utilities are all equal, and the approach is forced to behave as a *direct contact*, as there is no node better than the source to reach the destination. In order to show this, we measured the frequency of changes of the utilities associated to the nodes. To measure this index, as before we measure the number of changes in the utilities occurring in every time window of M minutes. The frequency drops to 0 after at most 6000 s. (the longest ICT) with DET, and a couple of days with RWP (Fig.3(a)).[2] When this occurs, all nodes are equal, and Δ_e becomes negative indicating that forwarding is through direct contact(Figg.1(a) and 3(b)).

By contrast, Greedy remembers the whole past history and properly detects the familiarity of a node with the destination. Thus, it perfectly catches the existence of 9 communities in DET (Fig.4(a)): for every node there is a sharp difference between the utilities of nodes outside and inside the destination community (Fig.4(b)). Hence, as soon as a node in the same community of the destination is encountered, it is adopted as a relay, and its utility guarantees that further forwarding can only occur within the destination community. After a short learning phase of around 2-3 hours, Greedy is able to characterize paths that reach all the destinations with performance comparable to that of $\mathcal{O}$.

In HCMM_det5, the challenge should be to capture the few travelers. Greedy utilities discriminate among travelers, nodes resident in a community different from that of a destination d, and nodes familiar with the destination. A sedentary node unfamiliar with d has utility for d equal to 0 (Fig.5(a), left part). By contrast, a node familiar with d encounters it many times and its counter is high (Fig.5(a), right part). Travelers show the behavior in Fig.5(b).[3] Travelers move around and encounter almost all nodes, but for few times: notice the difference in the x-axis between Fig.5(a) and (b). Fig.5(b) also shows that travelers do not visit all communities with uniform probability. The consequence on forwarding is that: (i) a node unfamiliar with d forwards the message to either a traveler or a node familiar with d that it happens to encounter; (ii) a traveler forwards the message to a node familiar with d. When a traveler is adopted as relay, its utility guarantees that further forwarding can only involve either a traveler more accustomed to visit the destination's group, or the destination's group itself. As soon as the destination's group is reached, forwarding is confined within it. Fig.5(c) shows that initially – during the first 8 hours roughly – counters (utilities) are not yet well differentiated, and direct contact is often used.

The behavior of Fresh and Prophet in HCMM_det5 can

[2]For the sake of comparison, the behavior with one setting where not all nodes encounter is shown. But, for readability, not all models are shown.

[3]Plots in Fig.5(a) and (b) are for two specific nodes, but their behavior is common to all residents and travelers respectively.

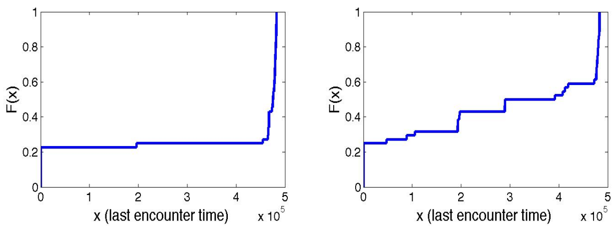

(a) (b)

Figure 7: Fresh in HCMM_det5: *ecdf*(utilities) for (a) a sedentary node, and (b) a traveler.

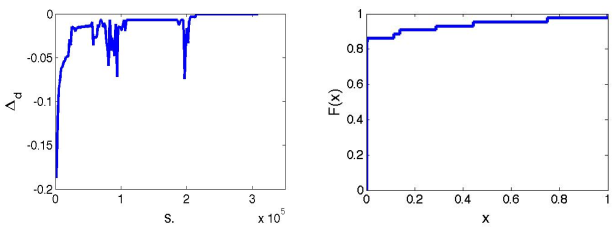

(a) (b)

Figure 8: HCMM_det5 – (a) Fresh: Δ_d vs. time. (b) Prophet: *ecdf*(utilities) for a node.

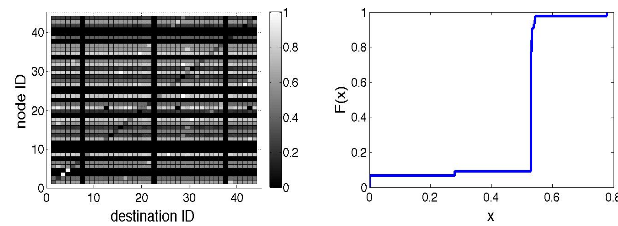

(a) (b)

Figure 9: SimBet in HCMM_det5: (a) utilities for every pair of nodes, and (b) *ecdf*(utilities) for a certain node.

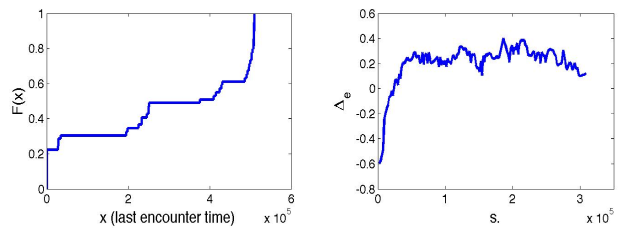

(a) (b)

Figure 10: Fresh in PMTR: (a) *ecdf*(utilities) for a node. (b) Δ_e vs. time.

be explained with similar arguments. Although the mobility pattern is close to DET, both the approaches perform better. In fact, despite Fresh has no memory, the habit of frequent encounters within groups continuously refreshes the time of the last encounter with a destination and lets emerge the familiarity.[4] In Fig.6, the value of the ICT's and the Fresh utilities are shown for all pairs of nodes, in (a) and (b) respectively. Sedentary nodes do not see unfamiliar nodes (Fig.7(a), left side), while have very recent encounters with nodes in their group (Fig.7(a), right side). By contrast, travelers have some nodes visited very recently (Fig.7(b), right side) – possibly of the last visited community – nodes visited in a more or less recent past (Fig.7(b), middle), and nodes never encountered (Fig.7(b), left side). The very same classification of nodes is achieved as with Greedy, with similar results. The successful learning of familiarities is brought into evidence by the variation of the reached destinations along time: after a learning period of around 7-8 hours, where some losses are experimented, all destinations are reached as with $\mathcal{O}$ (Fig.8(a)).

As far as Prophet is concerned, the aging mechanism virtually resets utilities deriving from sporadic visits to unfamiliar communities. As in Fresh, utilities of familiar nodes are continually refreshed. By contrast, the utility of travelers derives from two contrasting forces: on the one hand, the aging decreases the utility for nodes belonging to communities visited long ago. On the other hand, the transitivity mechanism increases the utility of travelers for nodes not encountered, but belonging to most visited communities (i.e., whose familiars have been encountered by the traveler). As a results, the utility values are slightly better differentiated (Fig.8(b)), and again acceptable performance is achieved.

In this environment, the approaches unable to single out useful relays are those partly or completely destination-independent, namely, Greedy Total and SimBet. With the former, a node continuously encountering its familiars – but never visiting unfamiliar groups – could have a high counter. With SimBet, each node has very homogeneous utilities for all destinations, indicating that the destination-independent component of the utility function is predominant (Fig.9(a)). Similarly, the *ecdf*'s over the utilities show only a big step (an example is shown in Fig.9(b)) distinguishing between the nodes never encountered and all the others. The nodes with high utilities are more than the travelers; they could also be nodes connecting two communities just because they belong to one and happened to visit the other. Such a node could be chosen as forwarder, but it gives no guarantees of encountering a destination resident in a different group.

The PMTR setting is quite similar to HCMM_det5, although the higher sedentariness and longer contact duration impose much higher latencies than in the other settings (Fig.1(b)). Yet, these characteristics help Fresh utilities in differentiating among familiar users (recently encountered), groups frequented less often, and users never seen (the multiple steps in Fig.10(a)). As in HCMM_det5, Fresh is able to learn these differences within roughly 12 hours (Fig.10(b)), and they are continuously refreshed thanks to users' habits in spite of the Fresh lack of memory. Noticeably, the differences among different degrees of familiarity are reported more sharply by Fresh utilities than by Greedy utilities (Fig. 11). We conjecture that this is due to the fact that the area is relatively small, and almost all pairs of nodes are likely to encounter soon or late. In such an environment, familiarity could be better caught by considering the duration of the encounters, rather than their number as Greedy does. For all nodes, the *ecdf* over Greedy utilities shows a behavior similar to that of travelers in HCMM_det5 (Fig.5(b)).

The other approaches are penalized by the characteristics of the environment. Greedy Total differentiates very well

[4]Notice the small average ICT in Table 1.

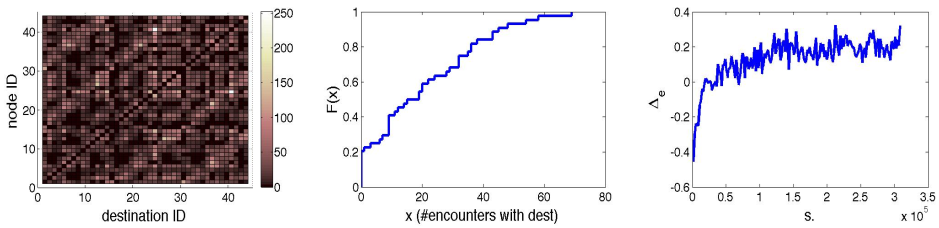

(*a*) (*b*) (*c*)

Figure 13: Greedy in HCMM_pro1: (*a*) utilities for every pair of nodes; (*b*) *ecdf*(utilities) for a node. (*c*) Δ_e.

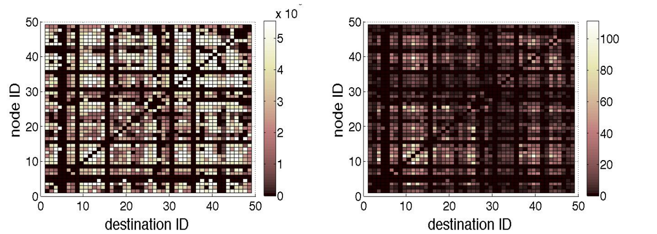

(*a*) (*b*)

Figure 11: PMTR: utilities for every pair of nodes for (*a*) Fresh, and (*b*) Greedy.

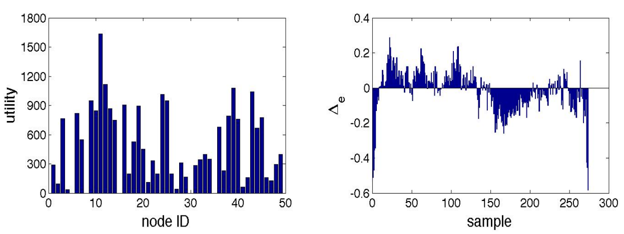

(*a*) (*b*)

Figure 12: PMTR – (*a*) Greedy Total: utilities. (*b*) Prophet: Δ_e vs. time.

among the popularity of different people (Fig.12(*a*)). We verified that nodes with high utilities correspond to individuals either with many cooperations, or whose office is in a central position where people frequently pass by. As before, utilities of SimBet mainly depend on betweenness. Though, as with HCMM_det5, destination-independence leads to inappropriate choice of the forwarders – which are not guaranteed to encounter the destination in the near future – and thus to poor performance. Due to high sedentariness, the utilities of Prophet are drastically reduced by the aging mechanism, making difficult for social structures to emerge, and thus hindering the identification of appropriate forwarders. As a consequence, direct contact is often used for message delivery (Fig.12(*b*)).

In HCMM_pro1, the accumulation of the whole history performed by Greedy makes it able to differentiate the communities, although mixed, as it emerges by both the utilities reported in Fig.13(*a*) (they are not uniform, as instead it happens in RWP) and the behavior of the *ecdf* over utilities (which shows more encounters with a subset of nodes) (Fig.13(*b*)). As a consequence, Greedy is able to discover the communities (although with a long learning phase: roughly 24 hours) and to move from a direct contact policy to an adequate choice of forwarders, as shown in Fig.13(*c*)

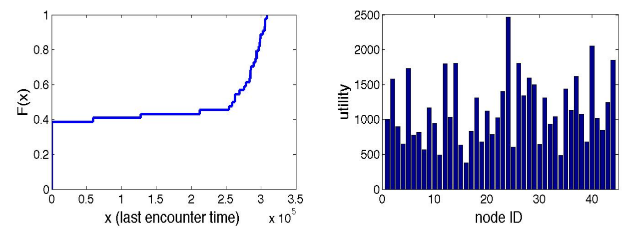

(*a*) (*b*)

Figure 14: HCMM_pro1 – (*a*) Fresh: *ecdf*(utilities) for a node. (*b*) Greedy Total: utilities of nodes.

The mixing of communities in HCMM_pro1 negatively affects Fresh ability of distinguishing different degrees of familiarity with a destination. Utilities either show nodes never encountered, or nodes recently encountered, which are all the others due to the mix (Fig.14(*a*)). Similarly, Prophet utilities do not make nodes habits to emerge. Mixing however helps the two approaches in reaching all the destinations, although with longer latencies than Greedy. By contrast, the destination-oblivious approach taken by both Greedy Total and SimBet, although able to differentiate among nodes with different popularity (Fig.14(*b*)), in spite of mixing communities, chooses forwarders that do not guarantee of eventually reaching the destinations, as highlighted by Fig.1(*c*).

In summary.

From the above considerations we can say that:

- in all environments, the maintenance of the whole history allows to accurately discriminate among relays;
- the approaches that either do not maintain or forget the past show good discrimination capability when the people habits continuously refresh the information about recurrent encounters, and there is a sharp difference between the encounter dynamics of familiar and unfamiliar nodes;
- as a consequence of the above argument, the approaches that maintain a qualitative record of the past encounters do not let recurrences to emerge, thus flattening the differences among relays;
- destination independency does not allow to identify the relays that are more likely approaching a given destination; this makes packets traveling away from the optimal paths.

5. CONCLUSIONS

In this paper, the behavior of five utility functions for message forwarding in ONs is studied, in five different mobility models. The results allow to characterize what mechanisms are able to discriminate among the usefulness of various relays depending on the environment characteristics. Hence, indications emerge that can drive the choice of the appropriate policy for computing the utilities according on the people mobility and the confinement of their communities.

Several developments are possible. Modifications of the mechanisms can be studied, in order to overcome difficulties that they may face in some environments. We are designing a utility mechanism able to adapt to the mobility characteristics of the people involved in an ON. As a future work, we plan to adapt some of the considered approaches in order they are able to "follow" several nodes pooled by a common interest, instead of a unicast destination.

6. ACKNOWLEDGMENT

The authors thank Chiara Boldrini for supplying the HCMM code. This work was funded partially by the Italian Ministry for Instruction, University and Research under the PRIN PEOPLENET (2009BZM837) Project, and partially by the European Commission under the FP7 SCAMPI (258414) Project.

7. REFERENCES

[1] S. Ahmed, S.S. Kanhere. HubCode: Message Forwarding using Hub-based Network Coding in Delay Tolerant Networks. In *Proc. MSWiM*, 2009, pp. 288-296.

[2] C. Boldrini, M. Conti, J. Jacopini, A. Passarella. HiBOp: a History Based Routing Protocol for Opportunistic Networks. In *Proc. WOWMOM 2007*, pp. 1-12.

[3] C. Boldrini, M. Conti, A. Passarella. Users Mobility Models for Opportunistic Networks: the Role of Physical Locations. In *Proc. IEEE WRECOM*, 2007, pp.1-6.

[4] –. BonnMotion: A mobility scenario generation and analysis tool. University of Bonn, Institute of Computer Science 4, Communication and Networked Systems. `http://net.cs.uni-bonn.de/wg/cs/applications/bonnmotion/`

[5] V. Borrel, F. Legendre, M. De Amorim, S. Fdida. SIMPS: Using Sociology for Personal Mobility. *IEEE/ACM Transactions on Networking*, 17(3), 2009, pp. 831-842.

[6] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. *Proc. IEEE INFOCOM 2006*, pp. 1-13.

[7] E. Daly, M. Haahr. Social Network Analysis for Routing in Disconnected Delay-Tolerant MANETs. In *Proc. MobiHoc 2007*, pp. 32-40.

[8] H. Dubois-Ferriere, M. Grossglauser, M. Vetterli. Age matters: efficient route discovery in mobile ad hoc networks using encounter ages. In *Proc. ACM MobiHoc 2003*, pp. 257-266.

[9] V. Erramilli, A. Chaintreau, M. Crovella, C. Diot. Diversity of forwarding paths in pocket switched networks. In *Proc. ACM/SIGCOMM IMC* (Oct. 2007), pp. 161-174.

[10] V. Erramilli, M. Crovella, A. Chaintreau, C. Diot. Delegation Forwarding. In *Proc. MobiHoc 2008*, pp. 251-259.

[11] S. Gaito, E. Pagani, G.P. Rossi. Strangers help friends to communicate in opportunistic networks. *Computer Network* 55(2), Elsevier, February 2011, 374-385.

[12] B. Han, P. Hui, V.S.A. Kumar, M.V. Marathe, G. Pei, A. Srinivasan. Cellular Traffic Offloading through Opportunistic Communications: A Case Study. *Proc. CHANTS 2010*.

[13] W. Hsu, T. Spyropoulos, K. Psounis, A. Helmy. Modeling Time-Variant User Mobility in Wireless Mobile Networks. In *Proc. IEEE INFOCOM 2007*, pp.758-766.

[14] P. Hui, J. Crowcroft, E. Yoneki. Bubble rap: social-based forwarding in delay tolerant networks. In *Proc. 9th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing*, 2008.

[15] K. Lee, S. Hong, S. J. Kim, I. Rhee, S. Chong. SLAW: A Mobility Model for Human Walks. In *Proc. IEEE INFOCOM 2009*.

[16] A. Lindgren, A. Doria, O. Schelńen. Probabilistic Routing in Intermittently Connected Networks. In *Proc. 1st International Workshop on Service Assurance with Partial and Intermittent Resources*, 2004, pp. 239-254.

[17] A. Mei, J. Stefa. SWIM: a Simple Model to Generate Small Mobile Worlds. In *Proc. IEEE INFOCOM 2009*.

[18] A. Mtibaa, M. May, M. Ammar, C. Diot. PeopleRank: Combining Social and Contact Information for Opportunistic Forwarding. In *Proc. IEEE Infocom 2010 Mini Conference*, pp. 1-5.

[19] M. Musolesi, C. Mascolo. Designing mobility models based on social network theory. *ACM SIGMOBILE Mobile Computing and Communications Review* 11(3), ACM Press (2007).

[20] L. Pelusi, A. Passarella, M. Conti. Opportunistic Networking: data forwarding in disconnected mobile ad hoc networks. *IEEE Communications Magazine* 44(11), November 2006, pp. 134-141.

[21] J. Su, A. Chin, A. Popivanova, A. Goel, E. de Lara. User mobility for opportunistic ad hoc networking. *Proceedings of IEEE WMCSA*, 2004.

[22] J. Whitbeck, M. Dias de Amorim, V. Conan. Plausible mobility: inferring movement from contacts. *Proceedings of ACM Workshop on Mobile Opportunistic Networking (MobiOpp)*, February 2010.

[23] J. Whitbeck, M. Dias de Amorim, Y. Lopez, J. Leguay, V. Conan. Relieving the Wireless Infrastructure: When Opportunistic Networks Meet Guaranteed Delays. *Proc. IEEE WoWMoM 2011*.