FOG: Fairness in Mobile Opportunistic Networking

Abderrahmen Mtibaa
School of Computer Science
Carnegie Mellon University
Doha, Qatar
Email: amtibaa@cmu.edu

Khaled A. Harras
School of Computer Science
Carnegie Mellon University
Doha, Qatar
Email: kharras@cs.cmu.edu

***Abstract*—**
The fundamental challenge in opportunistic networking, regardless of the application, is when and how to forward a message. Rank-based forwarding techniques currently represent one of the most promising methods for addressing this message forwarding challenge. While these techniques have demonstrated great efficiency in performance, they do not address the rising concern of fairness amongst various nodes in the network. Higher ranked nodes typically carry the largest burden in delivering messages, which creates a high potential of dissatisfaction amongst them. In this paper, we adopt a real-trace driven approach to study and analyze the tradeoff between the efficiency and fairness of rank-based forwarding techniques in mobile opportunistic networks. Our work comprises three major contributions. First, we quantitatively analyze the tradeoffs between fair and efficient environments. Second, we demonstrate how fairness coupled with efficiency can be achieved based on real mobility traces. Third, we propose FOG, a real-time distributed framework to ensure efficiency-fairness tradeoff using local information. Our data-driven experiment and analysis show that mobile opportunistic communication between users may fail with the absence of fairness in participating high-ranked nodes, and an absolute fair treatment of all users yields inefficient communication performance. Finally our analysis show that FOG ensures relative equality in the distribution of resource usage among neighbor nodes while keeping the success rate and cost performance near optimal.

I. Introduction

Mobile opportunistic networks interconnect nodes with heterogeneous contact rates, unpredictable mobility, and limited resources. These mobile nodes communicate relying on both the transport of messages as well as multi-hop forwarding. Current forwarding techniques [3], [26], [19], [21], [9] are generally designed to efficiently select the most likely relay nodes to deliver a message to its destination. Within those techniques, rank-based forwarding [5], [7], [8] represent one of the most promising methods for addressing this message forwarding challenge. This methods differ in the type of information used (*e.g.,* information acquired during contacts [7], [8], or social interaction between users [5], [24], [23]) as well as how it is used to rank nodes in the network. A node with a lower rank will forward messages to nodes with higher ranks. This solution, however, creates a high potential of dissatisfaction among high ranked nodes that carry a heavier burden compared to others. Providing fairness is therefore a crucial goal if such techniques are to be adopted in the future.

In mobile opportunistic networks, node cooperation is fundamental for the message delivery process. Therefore, the lack of node cooperation (*e.g.,* a node may refuse to act as a relay and settle for sending and receiving its own data) causes considerable delay degradation in the network. Fairness is then particularly important and challenging for mobile opportunistic networks and especially for rank-based forwarding algorithms since it acts as a major incentive for users to continue participating in the communication process. Previous studies have considered an absolute fair allocation of users' resources [1], [18], [20]. These studies have shown that while absolute fairness ensures a global sense of fairness among nodes, it deals with higher end-to-end delivery delays. It is then primordial to consider whether there exists a tradeoff relationship between fairness and efficiency. In this paper, we discuss whether there exists such a tradeoff, and propose FOG, a real-time distributed framework that enables existing rank-based forwarding algorithms to be both fair and efficient (*i.e.,* high message delivery success rate, and small end-to-end delay).

We adopt a real-trace driven approach to study and analyze the tradeoff between efficiency and fairness for rank-based forwarding techniques. Our work comprises three major contributions:

- First, we quantitatively analyze the impact of an absolute fair environment and an absolute efficient environment on the overall network performance. Our results show that in an unfair environment, selecting preferential, or popular, relay nodes in forwarding decisions is efficient and yields enhanced forwarding performance. In addition, these nodes are critical to the network operation. Mobile opportunistic communication may fail with the absence of the sense of fairness among these popular nodes. An absolute balancing of the across nodes, however, causes significant end-to-end delay and severe performance degradation (Section §III).
- Second, we discuss whether there exists a tradeoff relation between fairness and efficiency in opportunistic forwarding (Section §IV). Relying on our experimental datasets, we consider an offline approach to construct forwarding paths while ensuring both fairness and efficiency. We show that, while a tradeoff between fairness and efficiency may exist in a given dataset, a distributed

real-time approach is required in order for these paths to be efficiently discovered while utilizing only local information.
- Finally, we propose FOG, a real-time distributed framework that enables existing rank-based forwarding algorithms to be both fair and efficient. FOG relies on local information of a given node to ensure a fair distribution of the burden among its neighbor nodes. Our results show that our proposed framework ensures relative equality in the distribution of network resources while maintaining a high success rate, and near optimal cost; only 3% to 7% delivery rate performance regression, and 2% to 6% more message replicas in the network (Section §IV).

The remainder of this paper is organized as follows. Section II provides a brief overview of rank-based forwarding techniques and the experimental datasets used in our analysis. Section III discusses the Fairness vs. Efficiency tradeoff contributions. Our three contributions discussed earlier are then detailed in Sections III, IV, and V respectively. We then present the related work in the field of fair and efficient mobile opportunistic forwarding in Section VI. Finally, Section VII concludes the paper.

II. BACKGROUND & DATASETS

In this section we provide a brief overview on rank-based forwarding algorithms, and present the experimental datasets used in our analysis to improve the fairness of these forwarding algorithms.

A. Rank-based Forwarding Algorithms

Rank-based forwarding techniques represent one of the most promising methods for message forwarding in opportunistic networks. They differ in the type of information used, as well as how it is used, in order to rank nodes in the network. We distinguish between two type of rank-based forwarding techniques; contact-based ranking techniques, where information learned during contact, are used to rank nodes, and social-based ranking techniques, where social interactions between users are used to rank the nodes.

1) Social-based Ranking Algorithms: There are three well-known social-based ranking techniques that differ in the type of social metric used:

- Degree-Based Forwarding: Consists of forwarding messages to socially well connected nodes. Paths are then constructed according to a non-decreasing social node degree rule (more details in [22]).
- Centrality-Based Forwarding: The main idea behind this technique is that central nodes in social graphs are more likely to socialize with other people and therefore more likely to forward messages (more details in [5], [22]).
- PeopleRank: This is a distributed algorithm which ranks nodes in the social graph similar to what PageRank [2] does for web pages – *i.e.,* it measures the relative "importance" of a node in the social graph (more details in [23]).

2) Contact-based Ranking Algorithms: In the following, we present two of the most well-known contact-based ranking algorithms in the literature. As the name indicates, they use locally available contact information to rank nodes and decide whether to forward a message when two nodes meet. These two algorithms are:

- Last_Contact (LC): Node i forwards messages to node j if j has contacted any other node more recently than node i [7].
- Frequency (FR): Node i forwards the message to node j if j has more total contacts than node i [8].

B. Experimental Datasets

Our analysis is based on two datasets collected in conference environments [22], [4]. In addition to human mobility information, our datasets contain social relationships between the experimentalists. A summary of the corresponding parameters is provided in Table I.

	CoNext07	Infocom06
duration	3 days	3 days
mobility patterns	Bluetooth contacts	Bluetooth contacts
social patterns	MobiClique app. [22]	Facebook
# connected nodes	27	47
# edges	115	219
average degree	9.5	9.3
social diameter	4	4
median inter-contact	10 min	15 min
median contact time	240s	150s

TABLE I
CHARACTERISTICS OF OUR EXPERIMENTAL DATASETS

CoNext07: This dataset [22] contains mobility information and information about the social relationship between the participants during the ACM CoNext 2007 conference. During the experiment, the social networking application indicated when a contact, or a contact of a contact, was within Bluetooth range/neighborhood. This connection neighborhood was then displayed on the user's device which in turn could add new connections or delete existing connections based on the physical interaction consequent to the application notification.

Infocom06: This dataset [4] was collected with 78 participants during the IEEE Infocom 2006 conference. Experimentalists were asked to carry an experimental device (called iMote) with them at all times. These devices were logging all contacts between participating devices. Questionnaires were given to participants to fill theirs nationalities, languages, countries, cities, academic affiliations and topic of interests. Based on the forms filled by the experimentalists (*e.g.,* fist name, last name, interests, nationalities, affiliation, etc.), we consider, in this paper, a social graph based on users Facebook friendship graph (obtained offline and connecting all the experimentalists). Note that we did not use the implicit social information (*i.e.,* based on interests and affiliation) which is publically available in CRAWDAD[1].

[1]http://crawdad.cs.dartmouth.edu/meta.php?name=cambridge/haggle

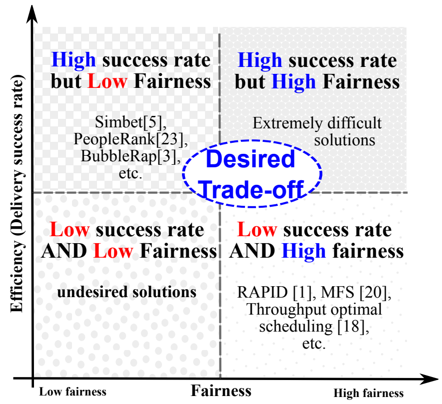

Fig. 1. Fairness-Efficiency tradeoff in mobile opportunistic communication

III. Fairness vs. Efficiency

Our discussion of efficiency and fairness highly depends on the definition of those two concepts. In this paper, we let "*efficiency*" denotes the successful delivery ratio of a given forwarding algorithm within t seconds. Higher efficiency means shorter end-to-end delay and more successful message delivery in the network.

To define fairness, one may consider it as a means to provide users with incentives to collaborate in mobile opportunistic communication – *i.e.,* encourage them to contribute to forwarding messages and remain longer in the network. We therefore define "*fairness*" as a relative equality in the distribution of resource usage among neighboring nodes in the network – *i.e.,* a forwarding algorithm is "fair" if the capacity assignment of a given node N in the network is equivalent to those of N's neighbors.

As summarized in Figure 1, while rank-based forwarding techniques have demonstrated great efficiency in performance [5], [24], [3], [23], they do not address the rising concern of fairness amongst various nodes in the network. Higher ranked nodes typically carry the largest burden in delivering messages, which creates a high potential of dissatisfaction amongst them. An absolute fair treatment of users, however, causes significant end-to-end delay and message delivery performance degradation [1], [18], [20]. Consequently, there is no technique to date that ensures both fairness and efficiency. It is therefore important to consider whether there exists a tradeoff relation between fairness and efficiency. In this paper, we answer the following question: Can we really fall in the desired trade-off region of Figure 1 (Section IV)?

A. Absolute Efficiency

As discussed earlier, efficiency deals with high potential of dissatisfaction among popular node. Let us assume that dissatisfied nodes decide to boycott the forwarding process. We show the impact that these popular nodes have on network performance when they refrain from message forwarding; popular node may be relevant but not critically needed for a network.

Methodology: We use a technique, which we call "*node removal*", to exclude nodes from the forwarding process. We study the impact of excluding popular nodes on the overall forwarding performance by applying this technique to the mobility traces we have. We also studied the impact of removing the unpopular nodes. From the sequence of delay-optimal paths in our datasets, we deduce the delay obtained by the optimal path at all time. We uniformly combine all the observations of a trace among all sources, destinations, and for every starting time (the time in seconds when the message m was generated by the source node S). We present this aggregated sample of observations via its empirical CDF. We plot the success rate of forwarding algorithms normalized by the success rate of flooding within a given message delivery delay. This technique is a relevant metric used to analyze forwarding algorithms and is discussed with more details in [4].

Figure 2 plots the empirical CDF of the normalized delivery success rate of the PeopleRank forwarding algorithm (an example of an efficient rank-based forwarding algorithm according to [23]). We have observed similar results for LC and FR algorithms. Figure 2-(a) shows the impact of excluding the most popular nodes on the successful delivery rate; 10% popular removal consists of removing the 10% most popular nodes from the original dataset (here, popularity rank is given by the PeopleRank values [23]). We show that mobile communication between users may fail with the absence of fairness, where excluding the most popular nodes from the forwarding process causes significant performance regress; if only 10% of the most popular nodes boycott the forwarding process, PeopleRank's success rate performance degrades by roughly 20% within a 10 minutes timescale.

However, PeopleRank performance shows insignificant regress when we exclude 10% of the most unpopular nodes; only 1.5% less within a 10 minutes timescale. This result is consistent with the fact that unpopular nodes do not contribute much to the forwarding process as the popular nodes do. Therefore, it is indeed shown that popular nodes in the network are more suitable than others to deliver a given message to its destination.

In Figure 2-(b), we plot the distribution of the successful delivery rate of PeopleRank as a function of different popular node removal percentages. As expected, excluding popular nodes from the forwarding process deteriorates the success rate, especially for small timescales. The probability to reach a destination within 10 minutes drops from 96% to 83% when 5% of the popular nodes are excluded. Moreover, the success rate regress is more severe when the most popular nodes are removed; while we show a 13% regress (from 96% to 83% within 10 minutes timescale) when only 5% of the most popular nodes are excluded, this regress is only 3% when we exclude 5% of the remaining 90% popular nodes (*i.e.,* we compare 10% and 15% removal nodes curves).

Figure 2-(c) shows that excluding unpopular nodes from the forwarding process does not seem to greatly impact the success rate of PeopleRank. At times, this removal may even increase

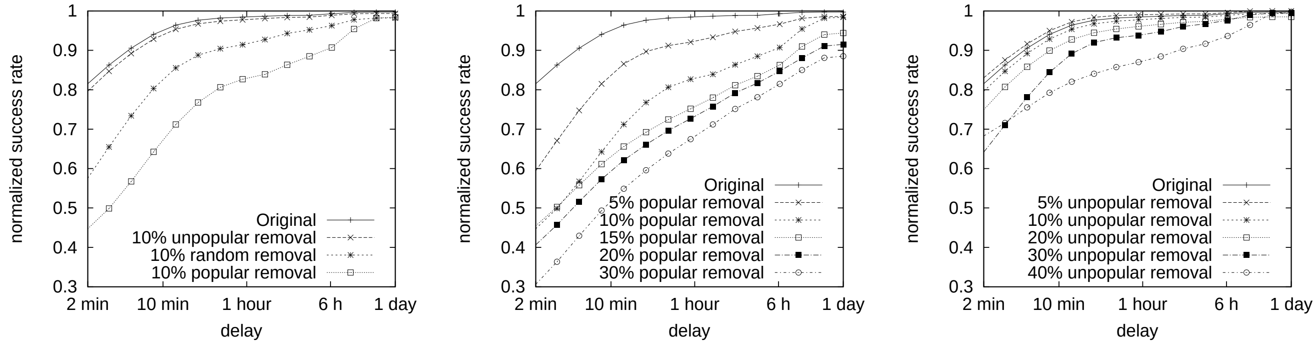

(a) 10% node removal (b) Popular node removal (c) Unpopular node removal

Fig. 2. *PeopleRank* performance with different node removal techniques (Infocom data set)

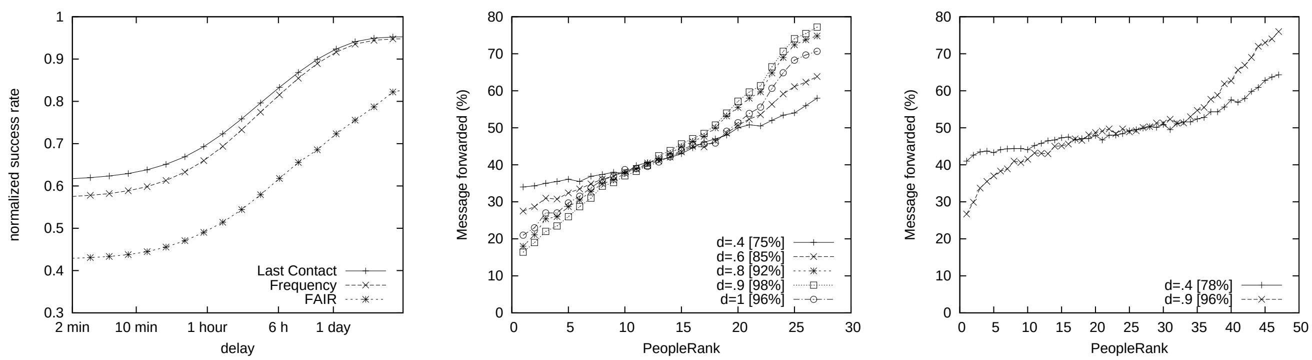

(a) Comparing rank-based algorithms (Last Contact, and Frequency) to FAIR success rate performance (using Infocom06 data set) (b) Normalized number of message forwarded and normalized success rate (in brackets) with regards of nodes' rank (CoNext07 dataset) (c) Normalized number of message forwarded and normalized success rate (in brackets) with regards of nodes' rank (Infocom06 dataset)

Fig. 3. Impact of fairness on the network performance

the performance when we remove only 5% of the most unpopular nodes. In contrast, removing the most unpopular nodes may, indeed, subtract high end-to-end delays from the overall distribution (*i.e.,* high delays are mainly caused by high waiting times to reach infrequently seen destinations).

As a summary, we observe that popular nodes in the network are more suitable than others to deliver a given message to its destination; selecting preferential relays in forwarding decisions gives better forwarding performance. It is indeed important to further satisfy popular nodes. Providing fairness is therefore crucial since the unfair treatment of users is considered as an disincentive to participation in the communication process.

B. Absolute Fairness

While fairness is our goal, absolute fairness amongst all nodes is not. In this section, we discuss the impact of absolute fairness on the overall network performance. Let us assume an absolute fair allocation of resources across nodes in the network. We perform an offline study of the path availability in our datasets, while each node forwards the same number of messages compared to all other nodes in the network. We ensure fairness by balancing cost across nodes. We call this offline fair path establishment technique, the FAIR algorithm. FAIR provides, by definition, a uniform fair allocation of user resources among nodes in the network.

Figure 3-(a) compares two contact-based ranking algorithms (*i.e.,* LC and FR) to the FAIR algorithm using the Infocom06 dataset. We show how an absolute fair allocation of resources leads to significant performance regression. The probability of success within 10 minutes decreases from 72% to 43% for Last Contact algorithm and from 60% to 43% for Frequency algorithm. The reason behind this regression is that an absolute fair balancing of cost across nodes causes significant end-to-end delay and success rate performance degradation.

PeopleRank [23] uses a simplistic technique in order to increase fairness across nodes. We study the impact of such techniques on the PeopleRank performance. PeopleRank allows us to tune the amount of the social information used through a damping factor d. [23] shows that social forwarding schemes are effective only for very accurate social information. The damping factor is then used to compensate the mismatch of the social interaction between users and their mobility patterns. Consequently, the more fairness we want, the closer to 0 we will chose the damping factor.

In Figure 3(b), we plot the number of message forwarded per PeopleRank node (nodes are ranked from the unpopular to the most popular node in the network) with regards to the damping factor d. This result is normalized by the number of messages forwarded if we run an epidemic forwarding algorithm. In addition, we indicate the normalized success rate (within 10 minutes) regarding each damping factor d in brackets on Figure 3(b)'s labels. We observe that the more unfair PeopleRank is (higher values of d), the better performance it can achieve; while PeopleRank algorithm is given the best distribution of resources among the nodes (less than 25% difference between the highest and the lowest used nodes in CoNext07 dataset when $d = 0.4$), it achieves only 75% of success rate performance within 10 minutes timescale. However, PeopleRank running with $d = 0.9$ gives 23% more success rate performance, and an important resources usage disproportion among the nodes (more than 55% difference). Similar results have been obtained in Figure 3(c) for the Infocom06 dataset. To summarize, an absolute fair balancing of cost across nodes causes significant end-to-end delay and success rate performance degradation.

IV. Can we be both Efficient & Fair?

In this section we quantitatively identify the tradeoff relationship between fairness and efficiency. We first provide some intuition regarding the desired fairness we seek while introducing our satisfaction index metric. We then discuss an offline approach to verify, based on our experimental datasets, the availability of paths that would satisfy both efficiency and fairness.

A. Desired Fairness & Satisfaction Index

In a mobile opportunistic network, we can roughly divide nodes into three different categories with respect to their popularity ranking; *popular*, *semi-popular* and *unpopular* nodes (as shown in Figure 4). While ranking nodes in the network may be considered to efficiently forward messages in DTN, it creates a high potential of dissatisfaction amongst nodes since they forward more messages compared to others (as shown by the solid line in Figure 4). On the other hand, an absolute fair allocation of resources yields poor forwarding performance (dashed line in Figure 4). Inspired by the previous results shown in Figure 2, we believe that the desired fairness is different from an absolute fair distribution of resources (optimal fairness). It is indeed important to increase the satisfaction of popular nodes by reducing their cost. Unpopular nodes, however, are unlikely impactful if they were involved in the forwarding process. Therefore, our goal is to further satisfy popular nodes by moving from a situation where popular nodes carry the largest burden in delivering messages to a "fair distribution" of this burden among *popular* and *semi-popular* nodes as shown by the red dotted line in Figure 4.

To reach this desired situation, we need a metric by which we can assess the level of fairness amongst the nodes. We define a *satisfaction index* (SI) as the difference between the message load distribution given by the forwarding process (current fairness) and uniform distribution among nodes

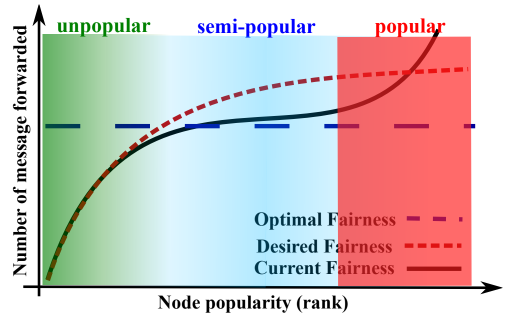

Fig. 4. Node categories with respect of their popularity rank

(desired fairness). Our goal is to construct paths between any source and destination that verify the condition below:

$$\exists C \geq 0, \forall i \in 1..n,$$
$$SI(N_i) = load(N_i) - \sum_{j=1}^{n} load(N_j) \geq 0 - C \quad (1)$$

where N_i represents a node i in the network, n is a total number of nodes, $load(N_i)$ is the current load at node N_i, and C is a positive constant integer. The goal is then to verify if there exists a relatively small positive integer C, such that we can construct paths between any source-destination pair and satisfy Equation 1.

B. Availability of Fair and Efficient Paths

We consider an offline approach to construct forwarding paths that ensures both fairness and efficiency based on global network information. In order to distribute the burden among semi-popular and popular nodes, we change the common forwarding rule used by most rank-based forwarding algorithms: Node A decides whether to forward a message m when it meets node B according to this inequality: $rank(B) > rank(A)$. This rule will exclude semi-popular nodes if m was generated by a high ranked node. We then apply the following rule: $rank(B) > rank(A) - \varepsilon$, where ε is a positive number, or a linear function. In this section, our goal is not to give a "magic" number or function of ε since this largely depends on the mobility characteristic present in the dataset. However, we discuss the fairness-efficiency tradeoff feasibility in a given dataset.

Figure 5 plots the normalized number of messages forwarded per node with respect to different values of ε. We show that when $\varepsilon = 0.2$, PeopleRank ensures a better distribution of the burden among popular and semi-popular nodes compared to the original forwarding rule (when $\varepsilon = 0$). On the other hand, the distribution of the number of messages forwarded by unpopular nodes seems rather insensitive to the variation of ε; it is indeed unlikely to meet these unpopular nodes, and the message will be delegated directly to popular and semi-popular nodes.

We now consider ε to be a function where increasing the nodes' rank r linearly increases ε value. We denote $\varepsilon(\alpha, \beta)$ as

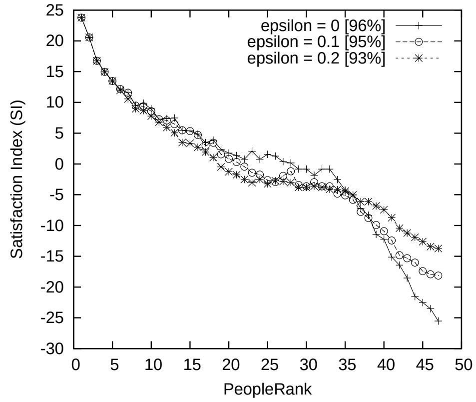

Fig. 5. Load distribution with respect to different ε values (Infocom06 data set)

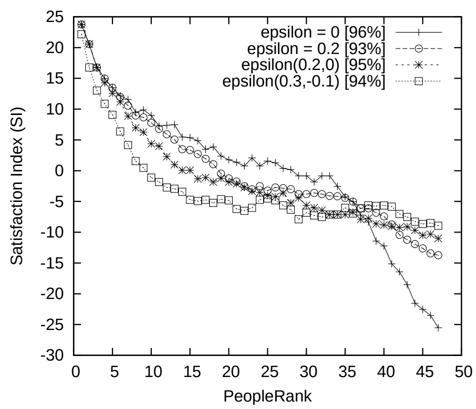

Fig. 6. Satisfaction factor with respect of different values of α and β (Infocom06 data set)

a function of the rank r, given by: $\varepsilon(\alpha, \beta) = \alpha * r + \beta$. The above forwarding rule can then be given by:

$$rank(B) > (1 - \alpha) * rank(A) + \beta \quad (2)$$

In Figure 6, we test different values of α and β, and plot the satisfaction index SI with respect to PeopleRank ranking. We show that the new rule given by Eq. 2 ensures a better distribution of the burden among the nodes and achieves a comparable success rate performance (only 1% or 2% less than optimal success rate performance). The new forwarding rule provides a fairly good satisfaction index performance as described in Equation 1 using a relatively small C constant ($C = 8$ for $\varepsilon(0.3, -0.1)$, and $C = 13$ for $\varepsilon(0.2, 0)$).

To summarize, our offline data driven approach shows that fair and efficient paths may exist in the network. However, it does not indicate whether theses paths can be found efficiently by a distributed algorithm.

V. A Real-time Distributed Framework for Fairness-Based Forwarding

We now propose a real-time distributed framework that helps rank-based forwarding algorithms utilize potential nodes to forward messages while satisfying the fairness property given by Equation 1. We call our systematic framework, *F*airness-based *O*pportunistic networkin*G* (*FOG*).

A. The FOG Framework

Whenever two nodes are within close proximity of each other, they separately run an update process. They first update their relative maximum burden max and relative minimum burden min per unit time. They then decide whether or not to forward a message m if the following condition is verified; (*i*) the encounter node N_j is the destination node of the message m, or (*ii*) N_j is higher ranked compared to N_i and its burden ($burden(N_j)$) did not exceed the relative average. Finally, the receiver of the message should update its actual burden value.

We note that, while the offline technique use different parameters to measure the satisfaction index SI among all nodes, FOG as a fully distributed framework, relies on local information at the node to estimate the distribution of the burden among neighboring nodes. We believe that nodes are satisfied by comparing themselves to their neighbors and acquaintances and require only relative equality amongst their neighbors in the network.

B. FOG Evaluation

In the following, we evaluate the performance of three state-of-the-art social forwarding algorithms (PeopleRank, LC, and FR) relying on the two previously described datasets (Infocom06, and CoNext07). We also use a data-driven artificial trace to evaluate the scalability of FOG in large scale social networks.

1) FOG in Small Social Environments: We adopt a real-trace driven approach to analyze and evaluate the tradeoff between the efficiency and fairness of FOG in order to provide a more realistic evaluation platform. We compare FOG's performance to the offline approach based on the real mobility traces shown in Table I.

In Figure 7, we compare the SI and the message distribution among nodes using FOG and the offline approach. In addition, we compare the success rate of these forwarding algorithms. We show that FOG outperforms all the offline rules we have tested and achieves one of the best SI/success rate tradeoffs. FOG also ensures a more fair distribution of the burden among the nodes; the SI of popular and semi-popular nodes remain close to zero and verify Equation 1 with a relatively small constant $C = 3$. FOG is explicitly prohibiting forwarding messages when the burden is not fairly distributed among neighbors; popular nodes are more likely to meet all other nodes as well as each other. The burden is therefore fairly distributed among them. We note that, FOG is also efficient given that it keeps the success rate close to the optimal (2% less compared to the original PeopleRank success rate performance).

2) Scalability of FOG: In order to study the scalability of our approach, we use a large scale artificial data-driven trace in our evaluation.

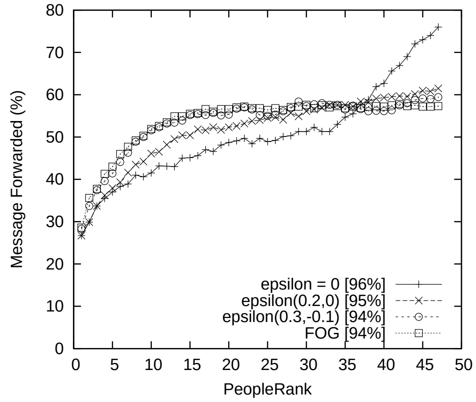

(a) Message forwarded

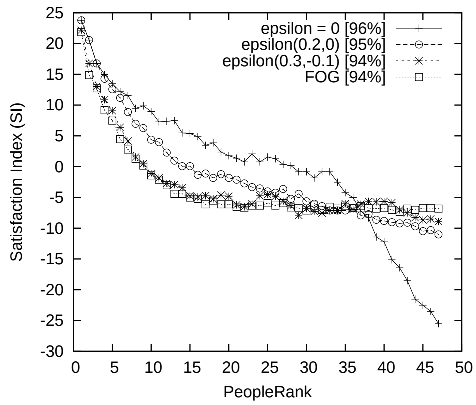

(b) SI

Fig. 7. Comparison of FOG and the offline technique (Infocom06 data set)

	San Francisco		
	CoNext07	Infocom06	Dartmouth
duration	3d	3d	3d
mobility pattern	Bluetooth	Bluetooth	WiFi
# nodes	27	47	100
median inter-contact	10mn	15mn	6mn
median contact time	240s	150s	160s

TABLE II
PROPERTIES OF THE SAN FRANSISCO MODIFIED DATA SET

The artificial data set used in our analysis is based on real human mobility information. We use the taxicabs San Francisco data set and incorporate traces for the duration of 3 days. Cabs interconnect different areas of San Francisco such as the airport, downtown, and the sunset area. We artificially add real human mobility patterns in each of the areas based on real traces (Infocom06, CoNext07, and Dartmouth campus; see Table II for more details). The resulting data set contains 3 communities interconnected by these cabs. Contacts between cabs and nodes within an area are added based on the same contact distribution of the corresponding area.

We evaluate the performance of FOG integrated with each of the three rank-based forwarding algorithms in large scale networks relying on the San Francisco modified dataset. Figure 8 plots for each rank-based algorithm (*a*) the normalized number of messages forwarded, and (*b*) the normalized success rate performance using the original algorithm and the FOG-based extension. We show that FOG-based extensions ensure better distribution of resources among the nodes while keeping the success rate performance close to optimal (original); we also show that FOG-FR (*i.e.,* FOG-based extension of FR algorithm) ensures fair allocation of the forwarding burden with a minimal decrease in success rate by only 4% within a 10 minutes timescale compared to the original FR success rate. Finally, we note that FOG-PeopleRank gives a significantly high improvement compared to the FAIR algorithm (described in Section III) in large scale networks: +20% in the success rate within a 10 minutes timescale. In PeopleRank, socially well ranked nodes carry the highest burden since they are more likely to forward a message to its destination. FOG-PeopleRank is therefore ensuring an acceptable tradeoff between efficiency and fairness in large scale networks.

3) Cost of FOG: Cost, in addition to end-to-end delay and success rate metrics, is a very important evaluation metric in opportunistic networks. We define the cost of a forwarding algorithm as the fraction of contacts involved in the forwarding process. FOG is designed to fairly distribute the burden among semi-popular and popular node. In this section, we study the impact of the fair distribution of the burden given by FOG on the overall cost.

Figure 9 compares the cost of FOG-PeopleRank to FAIR (absolute fairness) and the original version of PeopleRank using three different data sets (Infocom06, Conext07, and modified SF data sets). The original version of PeopleRank performs well and achieves the best normalized cost in all scenarios we have tested; PeopleRank, as a social-based ranking algorithm, was designed to provide a fairly well cost/efficiency tradeoff. We note that rank-based forwarding algorithms can be fair and keep the cost and the performance close to the optimal. FOG-PeopleRank outperforms the FAIR algorithm while using only 2% to 6% more replicas of the message m compared to the cost of the original PeopleRank algorithm.

It is crucial to be aware that, while the total number of replicas in the network increases, the cost increase per node is not significant as shown in the previous plots. Moreover, as the number of nodes in the network grows, reflecting more realistic deployment environments, the increase in cost becomes significantly smaller.

VI. Related Work

Most interactions between people rely on the establishment of the sense of fair treatment. Computer networking communication, and more particularly peer-to-peer file sharing applications and services take into consideration the fair treatment of users. Fairness is therefore particularly important and challenging since it is considered a major incentive for peer-to-peer service usage in today's Internet. Theses challenges are more critical in infrastructure-less wireless networks given the lack of centralized and trusted mechanisms that control the fair treatment of users. With the recent shift in research interest from centralized services to distributed services in mobile communication, fairness is becoming an interesting

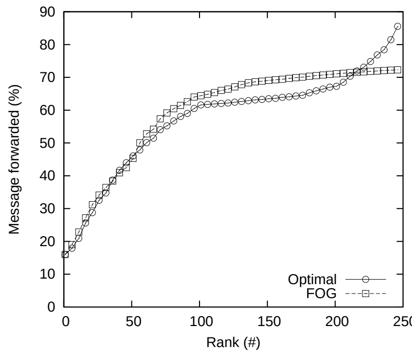

(a) LC algorithm

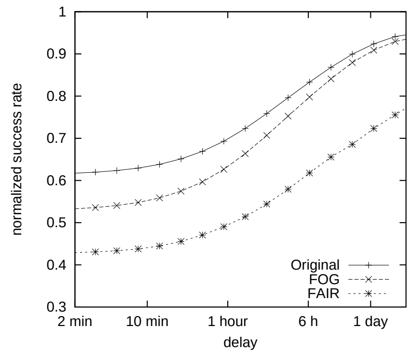

(b) LC algorithm

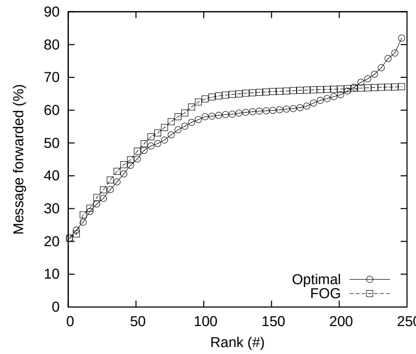

(c) FR algorithm

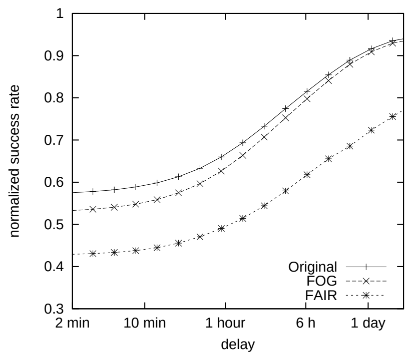

(d) FR algorithm

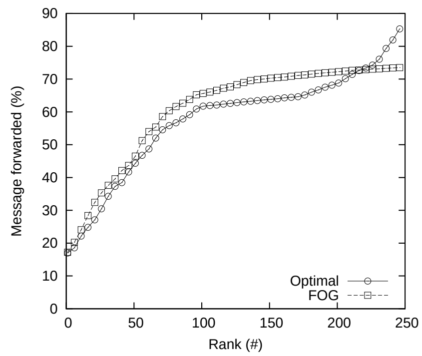

(e) PeopleRank algorithm

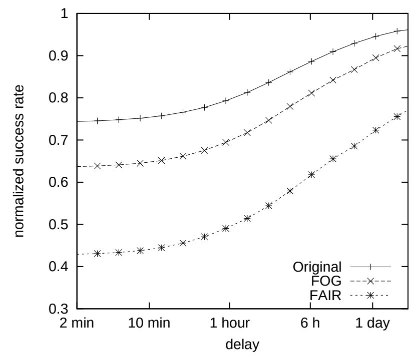

(f) PeopleRank algorithm

Fig. 8. Comparison of FOG and offline approach performance (using the modified SF dataset)

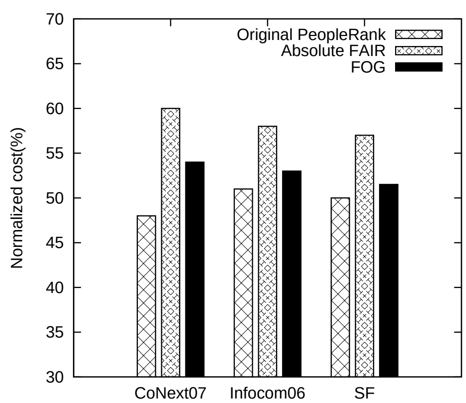

Fig. 9. Normalized cost of extended PeopleRank versions

field of investigation in many research topics such as resource allocation, congestion control, and network routing.

In peer-to-peer mobile networks such as mobile ad-hoc networks (MANETs) [6], [17], [11], [16], wireless sensor networks [25], [15], [12], or delay tolerant networks (DTNs) [14], [19], multi-hop wireless communication between users may fail with the absence of the sense of fairness between participating nodes. In DTNs, a device has to decide whether or not to forward data to an intermediate node that it encounters. Such forwarding decisions are typically guided by the desire to reduce the number of replicas of data items in the network to conserve bandwidth as well as by the desire to reduce end-to-end delay.

Current forwarding techniques in DTN are generally designed to efficiently, and excessively over-use popular nodes to guide and improve forwarding decisions [7], [8], [26], [1], [19], [21], [9], [27]. Rank-based forwarding techniques currently represent one of the most promising methods for addressing the message forwarding challenge [5], [23], [3]. Nodes in these techniques are ranked based on their social profiles or contact history to identify those that have a higher probability of successfully forwarding the message to the destination. While these techniques have demonstrated great efficiency in performance [5], [24], [23], [3], they do not address the rising concern of fairness amongst various nodes in the network. Higher ranked nodes typically carry the largest burden in delivering messages, which creates a high potential of dissatisfaction amongst them. Providing fairness is then an important networking goal since the unfair treatment of users is considered as a disincentive to participation in the communication process. It has been shown that ranking-based forwarding algorithms provide good performance relying on identifying and overusing popular nodes in the network [5], [24], [23]. Such well ranked nodes are more likely than others to deliver a message to its destination in a shorter delay. As a direct consequence, an absolute fair treatment of users causes a significant end-to-end delay and message delivery performance degradation. It is then primordial to consider whether there exists a tradeoff relationship between fairness and efficiency.

Fair sharing of resources has been largely studied in the context of classical networks. Previous work has been inspired by the well known max-min fairness [28] or Jain's fairness index [13] in order to improve a fair allocation/scheduling of resources in the Internet. In the context of wireless networks, researchers use the link quality variation of access points to maximize aggregate throughput [18], [20] . In the context of DTNs many assumptions are made regarding resource constraints (storage and bandwidth), and strictly adhere to max-min fairness for end-to-end delay minimization [10], [1]. In this paper, we address the above issue more generally. We assume unlimited resources in the network and propose a real time distributed framework to improve forwarding decisions in order to avoid dissatisfaction among popular nodes, and therefore ensure an efficiency-fairness tradeoff using local information.

VII. Conclusion

Our work addresses the rising concern of fairness amongst various nodes in mobile-based opportunistic networks. Rank-based forwarding algorithms are typically designed to reduce the number of data replicas in the network to conserve bandwidth and reduce end-to-end delay. In these algorithms, higher ranked nodes carry a much heavier burden in delivering messages, which can create high levels of dissatisfaction amongst them. In this paper, we have adopted a real-trace driven approach to study and analyze the tradeoff between the efficiency and fairness of rank-based forwarding algorithms in mobile opportunistic networks. Our preliminary results show that mobile opportunistic communication between users may fail with the absence of the perception of fairness. Furthermore, an absolute fair treatment of users also yields inefficient communication performance. Our offline approach, which relies on real mobility traces, shows that a tradeoff between fairness and efficiency can be considered in the forwarding process. We have proposed a real-time distributed framework called FOG to ensure efficiency-fairness tradeoff using local information within a node's neighborhood. Our analysis shows that FOG ensures relative equality in the distribution of resource usage among neighbor nodes while maintaining a high delivery success rate, and cost performance close to optimal.

References

[1] A. Balasubramanian, B. Levine, and A. Venkataramani. Dtn routing as a resource allocation problem. *SIGCOMM Comput. Commun. Rev.*, 37(4):373–384, 2007.

[2] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Syst.*, 30(1-7):107–117, 1998.

[3] J. Burgess, B. Gallagher, D. Jensen, and B. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networking. In *Proceedings of IEEE Infocom*, 2006.

[4] A. Chaintreau, A. Mtibaa, L. Massoulié, and C. Diot. The diameter of opportunistic mobile networks. In *Proceedings of ACM CoNext*, 2007.

[5] E. M. Daly and M. Haahr. Social network analysis for routing in disconnected delay-tolerant manets. In *MobiHoc '07: Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*, pages 32–40, New York, NY, USA, 2007. ACM.

[6] S. R. Das. ad hoc on-demand distance vector (aodv) routing. 2003.

[7] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli. Age matters: efficient route discovery in mobile ad hoc networks using encounter ages. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 257–266, 2003.

[8] V. Erramilli, A. Chaintreau, M. Crovella, and C. Diot. Diversity of forwarding paths in pocket switched networks. In *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 161–174, New York, NY, USA, 2007. ACM.

[9] V. Erramilli, M. Crovella, A. Chaintreau, and C. Diot. Delegation forwarding. In *MobiHoc '08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pages 251–260, New York, NY, USA, 2008. ACM.

[10] S. Guo and S. Keshav. Fair and efficient scheduling in data ferrying networks. In *Proceedings of the 2007 ACM CoNEXT conference*, CoNEXT '07, pages 13:1–13:12, New York, NY, USA, 2007. ACM.

[11] Z. J. Haas, J. Deng, B. Liang, P. Papadimitratos, and S. Sajama. Wireless ad hoc networks, 2002.

[12] W. Henzeilman, J. Kulik, and H. Balaksrishnan. Adaptive protocols for information dissemination in wireless sensor networks. In *Proceedings of ACM Mobicom*, 1999.

[13] R. Jain, D. Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems, 1998.

[14] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proceedings of ACM SIGCOMM*, 2004.

[15] S. Jain, R. C. Shah, G. Borriello, W. Brunette, and S. Roy. Exploiting mobility for energy efficient data collection in sensor networks. In *WiOpt*, March 2004.

[16] D. Johnson and A. Lysko. Comparison of manet routing protocols using a scaled indoor wireless grid. *Mob. Netw. Appl.*, 13(1-2):82–96, 2008.

[17] D. B. Johnson, D. A. Maltz, and J. Broch. Dsr: The dynamic source routing protocol for multi-hop wireless ad hoc networks. In *In Ad Hoc Networking, edited by Charles E. Perkins, Chapter 5*, pages 139–172. Addison-Wesley, 2001.

[18] S. S. Kulkarni and C. Rosenberg. Opportunistic scheduling for wireless systems with multiple interfaces and multiple constraints. In *Proceedings of the 6th ACM international workshop on Modeling analysis and simulation of wireless and mobile systems*, MSWIM '03, pages 11–19, New York, NY, USA, 2003. ACM.

[19] A. Lindgren, A. Doria, and O. Schelén. Probabilistic routing in intermittently connected networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 7(3):19–20, 2003.

[20] Y. Liu and E. Knightly. Opportunistic fair scheduling over multiple wireless channels. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, 2003.

[21] S. Merugu, M. Ammar, and E. Zegura. Routing in space and time in networks with predictable mobility. Technical report, 2004.

[22] A. Mtibaa, A. Chaintreau, J. LeBrun, E. Oliver, A.-K. Pietilainen, and C. Diot. Are you moved by your social network application? In *WOSN'08: Proceedings of the first workshop on Online social networks*, pages 67–72, New York, NY, USA, 2008. ACM.

[23] A. Mtibaa, M. May, M. Ammar, and C. Diot. Peoplerank: Social opportunistic forwarding. In *Proceedings of IEEE INFOCOM (mini conference)*. IEEE, 2010.

[24] A. Mtibaa, M. May, and M. H. Ammar. On the relevance of social information to opportunistic forwarding. In *MASCOTS'10*, pages 141–150, 2010.

[25] S. J. R. C. Shah, S. Roy and W. Brunette. Datamules: Modelling a three tiered architecture for sparse sensor networks. In *Proceedings of the First IEEE International Workshop on Sensor Network Protocols and Applications (SNPA)*, 2003.

[26] N. Sarafijanovic-Djukic, N. S. Janovic-djukic, and M. Grossglauser. Last encounter routing under random waypoint mobility. In *In Networking*, pages 974–988, 2004.

[27] C. Sengul, A. Carneiro Viana, R. Friedman, M. Bertier, and A.-M. Kermarrec. Adaptive Forwarding to Match Mobility Characteristics in Delay Tolerant Networks. Research Report RR-6816, INRIA, 2009.

[28] H. Zhang. Service disciplines for guaranteed performance service in packet-switching networks. *Proceedings of the IEEE*, 83(10):1374 – 1396, Oct. 1995.