

BUBBLE Rap: Social-Based Forwarding in Delay-Tolerant Networks

Pan Hui, *Member, IEEE*, Jon Crowcroft, *Fellow, IEEE*, and Eiko Yoneki, *Member, IEEE*

Abstract—The increasing penetration of smart devices with networking capability form novel networks. Such networks, also referred as pocket switched networks (PSNs), are intermittently connected and represent a paradigm shift of forwarding data in an ad hoc manner. The social structure and interaction of users of such devices dictate the performance of routing protocols in PSNs. To that end, social information is an essential metric for designing forwarding algorithms for such types of networks. Previous methods relied on building and updating routing tables to cope with dynamic network conditions. On the downside, it has been shown that such approaches end up being cost ineffective due to the partial capture of the transient network behavior. A more promising approach would be to capture the intrinsic characteristics of such networks and utilize them in the design of routing algorithms. In this paper, we exploit two social and structural metrics, namely *centrality* and *community*, using real human mobility traces. The contributions of this paper are two-fold. First, we design and evaluate BUBBLE, a novel social-based forwarding algorithm, that utilizes the aforementioned metrics to enhance delivery performance. Second, we empirically show that BUBBLE can substantially improve forwarding performance compared to a number of previously proposed algorithms including the benchmarking history-based PROPHET algorithm, and social-based forwarding SimBet algorithm.

Index Terms—Social networks, forwarding algorithms, delay-tolerant networks, pocket-switched networks, centrality, community detection.



1 INTRODUCTION

WE envision a future in which a multitude of devices carried by people are dynamically networked. We aim to build pocket switched networks (PSN) [1], a type of delay-tolerant networks (DTN) [2] for such environments. A PSN utilizes contact opportunities to allow humans to communicate without network infrastructure. We propose an efficient data forwarding mechanism over time evolving graphs of the PSN [3], that copes with dynamical, repeated disconnection, and rewiring. With such scenarios, end-to-end delivery through traditional routing algorithms is rarely applicable.

Many MANETs and some DTN routing algorithms [4], [5] provide forwarding by building and updating routing tables whenever mobility occurs. We believe this approach is not appropriate for a PSN, since mobility is often unpredictable and topology structure is highly dynamic. Rather than exchange much control traffic to create unreliable routing structures, which may only capture the “noise” of the network, we prefer to search for some characteristics of the network, which are less volatile than mobility. A PSN is formed by people. Hence, social metrics are intrinsic properties to guide data forwarding in such kinds of human networks. Furthermore, if we can detect these social mobility patterns online in a decentralized way, we can apply these algorithms in practice.

In this paper, we focus on two key social metrics: community and centrality. Cooperation binds, but also divides human society into communities. For an ecological community, the idea of correlated interaction means that an organism of a given type is more likely to interact with another organism of the same type than with a randomly chosen member of the population [6]. This correlated interaction concept also applies to human, so we can exploit this kind of community information to select forwarding paths. Within a community, some people are more popular, and interact with more people than others (i.e., have high *centrality*); we call them hubs. In this paper, we will exploit community and centrality for data forwarding in PSNs.

Methodologically, community detection [7], [8] can help us to uncover and understand local community structure in both offline mobile trace analysis and online applications, and is therefore helpful in designing good strategies for information dissemination. Freeman [9] defined several centrality metrics to measure the importance of a node in a network. Betweenness centrality measures the number of times a node falls on the shortest path between two other nodes. This concept is also valid in a DTN. In a PSN, it can represent the importance of a node as a potential traffic relay for other nodes in the system. The main contributions of this paper are to answer following questions:

- P. Hui is with Deutsche Telekom Laboratories, Ernst-Reuter-Platz 7, Berlin 10587, Germany. E-mail: pan.hui@telekom.de.
- J. Crowcroft and E. Yoneki are with the Computer Laboratory, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge CB3 0FD, UK. E-mail: {Jon.Crowcroft, Eiko.Yoneki}@cl.cam.ac.uk.

Manuscript received 25 May 2009; revised 12 Oct. 2010; accepted 25 Oct. 2010; published online 20 Dec. 2010.

For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-2009-05-0189. Digital Object Identifier no. 10.1109/TMC.2010.246.

1. How does the variation in node popularity help us to forward in a PSN?
2. Are communities of nodes detectable in PSN traces?
3. How well does social-based forwarding work, and how does it compare to other forwarding schemes in a real (emulated) environment?
4. Can we devise a fully decentralized way for such schemes to operate?

TABLE 1
Characteristics of the Five Experimental Data Sets

Experimental data set	Infocom05	Hong-Kong	Cambridge	Infocom06	Reality
Device	iMote	iMote	iMote	iMote	Phone
Network type	Bluetooth	Bluetooth	Bluetooth	Bluetooth	Bluetooth
Duration (days)	3	5	11	3	246
Granularity (seconds)	120	120	600	120	300
Number of Experimental Devices	41	37	54	98	97
Number of internal contacts	22,459	560	10,873	191,336	54,667
Average # Contacts/pair/day	4.6	0.084	0.345	6.7	0.024
Number of external devices	264	868	11,357	14,036	NA
Number of external contacts	1,173	2,507	30,714	63,244	NA

To preview our answers to the above questions, we evaluate the impact of community and centrality on forwarding, and propose BUBBLE, a hybrid algorithm, that selects high centrality nodes and community members of destination as relays. We demonstrate a significant improvement in forwarding efficiency over a number of previously proposed state-of-the-arts algorithms including PROPHET algorithm [5], which uses patterns of movement, rather than the longer term social relationships on which our BUBBLE scheme rests, and the social-based forwarding SimBet algorithm [10]. In a PSN, there may be no a priori information. By definition, we are also in a decentralized world without access to infrastructure. Therefore, distributed detection and dissemination of node popularity and node communities, and the use of these for forwarding decisions are crucial. We verify that this is not only possible, but works well in terms of message delivery performance and efficiency compared to prior schemes.

The rest of this paper is structured as follows: We introduce the experimental data sets in Section 2, describe the contact graphs and inferring human communities in Section 3. In Section 4, we examine the heterogeneity in centrality. We show evaluation methodology in Section 5, and results including discussions from Section 6 to 9, followed by the related works in Section 10. Finally, we conclude the paper with a brief summary of contributions in Section 11.

2 EXPERIMENTAL DATA SETS

In this paper, we use four experimental data sets gathered by the Hagggle Project¹ over two years, referred to as *Infocom05*, *HongKong*, *Cambridge*, *Infocom06*; one data set from the MIT Reality Mining Project [11], referred to as *Reality*. Previously, the characteristics of these data sets, such as intercontact and contact distribution, have been explored in several studies [12], [13], [14], to which we refer the reader for further background information. These five data sets cover a rich diversity of environments, ranging from busy metropolitan city (*HongKong*) to quiet university town (*Cambridge*), with an experimental period from several days (*Infocom06*) to almost one year (*Reality*).

- In *Infocom05*, the devices were distributed to approximately 50 students attending the Infocom student workshop. Participants belong to different social communities (depending on their country of origin, research topic, etc.).

- In *Hong-Kong*, the people carrying the wireless devices were chosen independently in a Hong-Kong bar, to avoid any particular social relationship between them. These people have been invited to come back to the same bar after a week. They are unlikely to see each other during the experiment.
- In *Cambridge*, the iMotes were distributed mainly to two groups of students from the University of Cambridge Computer Laboratory, specifically undergraduate year1 and year2 students, and also some PhD and masters' students. This data set covers 11 days.
- In *Infocom06*, the scenario was very similar to *Infocom05* except that the scale is larger, with 80 participants. Participants were selected so that 34 out of 80 form 4 subgroups by academic affiliations.
- In *Reality*, 100 smart phones were deployed to students and staff at MIT over a period of nine months. These phones were running software that logged contacts with other Bluetooth-enabled devices by doing Bluetooth device discovery every five minutes.

The five experiments are summarised in Table 1. A remark about the data sets is that the experiments do not have the same granularity and the finest granularity is limited to 120 seconds. This is because of the trade-off between the duration of the experiments and the accuracy of the samplings.

3 INFERRING HUMAN COMMUNITIES

In PSN, the social network could map to the computer network since people carry the computing devices. To answer the question whether communities of nodes are detectable in PSN traces we need community detection algorithms. In this section, we introduce and evaluate two centralized community detection algorithms: *K-CLIQUE* by Palla et al. [15] and weighted network analysis (WNA) by Newman [16]. We use these two centralized algorithm to uncover the community structures in the mobile traces. We believe our evaluation of these algorithms can be useful for future traces gathered by the research community.

Many centralized community detection methods have been proposed and examined in the literature (see the recent review papers by Newman [7] and Danon et al. [8]). The criteria we use to select a centralized detection method are the ability to uncover overlapping communities, and a high degree of automation (low manual involvement). In real

1. <http://www.hagggleproject.org>.




Fig. 1. The distribution of pair-wise contact durations.

human societies, one person may belong to multiple communities and, hence, it is important to be able to detect this feature. The K -CLIQUE method satisfies this requirement, but was designed for binary graphs, thus we must threshold the edges of the contact graphs in our mobility traces to use this method, and it is difficult to choose an optimum threshold manually [15]. On the other hand, WNA can work on weighted graphs directly, and does not need thresholding, but it cannot detect overlapping communities [16]. Thus, we chose to use both K -CLIQUE and WNA; they each have useful features that complement one another.

3.1 Contact Graphs

In order to help us to present the mobility traces and make it easier for further processing, we introduce the notion of *contact graph*. The way we convert human mobility traces into weighted contact graphs is based on the number of contacts and the contact duration, although we could use other metrics. The nodes of the graphs are the physical nodes from the traces, the edges are the contacts, and the weights of the edges are the values based on the metrics specified such as the number of contacts during the experiment. We can measure the relationship between two people by how many times they meet each other, and how long they stay with each other. We naturally think that if two people spend more time together or see each other more often, they are in a closer relationship.

First, we find the distribution of contact durations and number of contacts for the two conference scenarios are quite similar. To prevent redundancy, in the later sections we only selectively show one example, in most cases *Infocom06*, since it contains more participants.

Figs. 1 and 2 show the contact duration and number of contacts distribution for each pair in four experiments. For the *HongKong* experiment, we include the external device because of the network sparseness, but for the other three experiments we use only the internal devices. These contact graphs created are used for the community detection in the following sections.

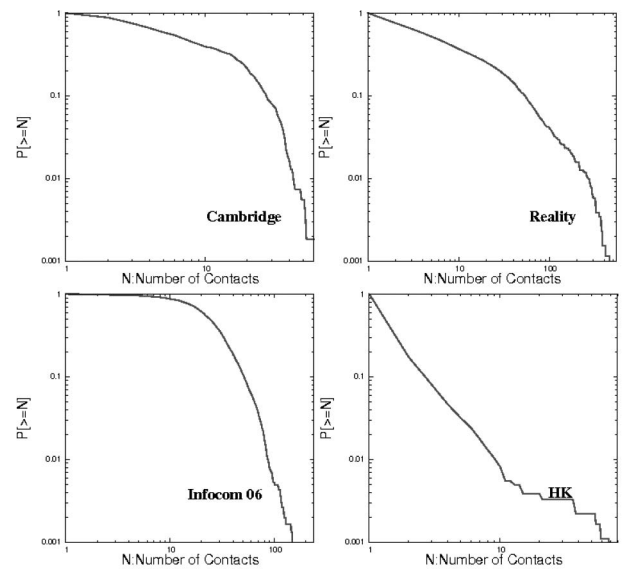


Fig. 2. The distribution of pair-wise number of contacts.

3.2 K-CLIQUE Community Detection

Palla et al. [15] define a k -clique community as a union of all k -cliques (complete subgraphs of size k) that can be reached from each other through a series of adjacent k -cliques, where two k -cliques are said to be adjacent if they share $k - 1$ nodes. As k is increased, the k -clique communities shrink, but on the other hand become more cohesive since their member nodes have to be part of at least one k -clique. We have applied this on all the data sets above. Fig. 3 shows the 3-clique communities in the *Infocom06* data set. More detailed descriptions about the k -clique communities on these data sets can be found in our previous work [17], [18].

3.3 Weighted Network Analysis

In this section, we implement and apply Newman's weighted network analysis (WNA) for our data analysis [16]. This is an extension of the unweighted *modularity* proposed in [19] to a weighted version, and use this as a measurement of the fitness of the communities it detects.

For each community partitioning of a network, one can compute the corresponding modularity value using the following definition of *modularity* (Q):

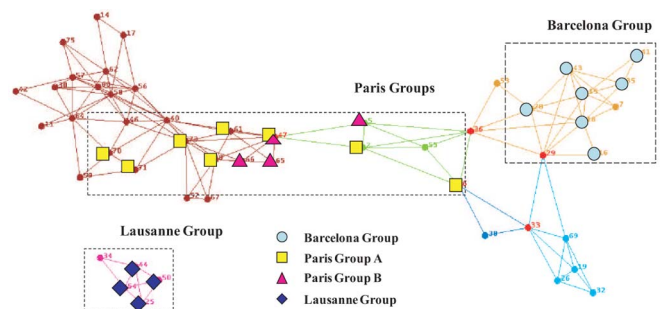


Fig. 3. 3-clique communities based on contact durations with weight threshold that equals 20,000 s (*Infocom06*; circles, Barcelona group; squares, Paris group A; triangles, Paris group B; diamonds, Lausanne group).

TABLE 2
Communities Detected from the Four Data Sets

Dataset	Info06	Camb	Reality	HK
Q_{max}	0.2280	0.4227	0.5682	0.6439
Max. Community Size	13	18	23	139
No. Communities	4	2	8	19
Avg. Community Size	8.000	16.500	9.875	45.684
No. Community Nodes	32	33	73	868
Total No. of Nodes	78	36	97	868

$$Q = \sum_{vw} \left[\frac{A_{vw}}{2m} - \frac{k_v k_w}{(2m)^2} \right] \delta(c_v, c_w), \quad (1)$$

where A_{vw} is the value of the weight of the edge between vertices v and w , if such an edge exists, and 0 otherwise; the δ -function $\delta(i, j)$ is 1 if $i = j$ and 0 otherwise; $m = \frac{1}{2} \sum_{vw} A_{vw}$; k_v is the degree of vertex v defined as $\sum_w A_{vw}$; and c_i denotes the community vertex i belongs to. *Modularity* is defined as the difference between this fraction and, the fraction of the edges that would be expected to fall within the communities if the edges were assigned randomly, but we keep the degrees of the vertices unchanged. The algorithm is essentially a genetic algorithm, using the modularity as the measurement of fitness. Rather than selecting and mutating current best solutions, we enumerate all possible merges of any two communities in the current solution, and evaluate the relative fitness of the resulting merges, and choose the best solution as the seed for the next iteration.

Table 2 summarizes the communities detected by applying WNA on the four data sets. According to Newman [16], nonzero Q values indicate deviations from randomness; values around 0.3 or more usually indicate good divisions. For the *Infocom06* case, the Q_{max} value is low; this indicates that the community partition is not very good in this case. This also agrees with the fact that in a conference the community boundary becomes blurred. For the *Reality* case, the Q value is high; this reflects the more diverse campus environment. For the *Cambridge* data, the two groups spawned by WNA is exactly matched the two groups (1st year and 2nd year) of students selected for the experiment.

These centralized community detection algorithms give us rich information about the human social clustering and are useful for offline data analysis on mobility traces collected. We can use them to explore structures in the data and, hence, design useful forwarding strategies, security measures, and killer applications.

4 HETEROGENEITY IN CENTRALITY

In human society, people have different levels of popularity: salesmen and politicians meet customers frequently, whereas computer scientists may only meet a few of their colleagues once a year [17]. Here, we want to employ heterogeneity in popularity to help design more efficient forwarding strategies: we prefer to choose popular hubs as relays rather than unpopular ones. SimBet routing algorithm [10] also uses the concept of centrality for choosing relays. We will compare the performance of both algorithms in Section 8.




Fig. 4. Number of times a node as relays for others on four data sets.

A temporal network or time evolving network is a kind of weighted network. The centrality measure in traditional weighted networks may not work here, since the edges are not necessarily concurrent (i.e., the network is dynamic and edges are time-dependent). Hence, we need a different way to calculate the centrality of each node in the system. Our approach is as follows:

1. Carry out a large number of emulations of unlimited flooding with different uniformly distributed traffic patterns created using the *HaggleSim* emulator (Section 5.1).
2. Count the number of times a node acts as a relay for other nodes on all the shortest delay deliveries. Here the shortest delay delivery refers to the case when the same message is delivered to the destination through different paths, where we only count the delivery with the shortest delay.

We call the number calculated above the *betweenness centrality* of this node in this temporal graph. Of course, it can be normalized to the highest value found. Here we use unlimited flooding since it can explore the largest range of delivery alternatives with the shortest delay. This definition captures the spirit of Freeman centrality [9].

Initially, we only consider a homogeneous communication pattern, with equal likelihood of every destination, and we do not weight the traffic matrix by locality. We then calculate the global centrality value for the whole homogeneous system. Later, we will analyse the heterogeneous system (Section 8).

Fig. 4 shows the number of times a node falls on the shortest paths between all other node pairs. We can treat this simply as the centrality of a node in the system. We observe very wide heterogeneity in each experiment. This clearly shows that there is a small number of nodes which have extremely high relaying ability, and a large number of nodes that have moderate or low centrality values, across all experiments. One interesting point from the HK data is that the node showing highest delivery power in the figure

TABLE 3
Statistics about Normalized Node
Centrality in Four Experiments

Experimental dataset	30th percentile	Mean	70th percentile
Cambridge	0.052	0.220	0.194
Reality	0.005	0.070	0.050
Infocom06	0.121	0.188	0.221
Hong Kong	0.000	0.017	0.000

is actually an external node. This node could be some popular hub for the whole city, i.e., postman or a newspaper man in a popular underground station, which relayed a certain amount of cross city traffic. The 30th, 70th percentiles, and the means of normalized individual node centrality are shown in Table 3. These numbers summarize the statistical property of the centrality values for each system shown in Fig. 4.

5 INTERACTION AND FORWARDING

In the first half of this paper, we have shown the existence of heterogeneity at the level of individuals and groups, in all the mobility traces. This motivates us to consider a new heterogeneous model of human interaction and mobility.

Cliques and Community. We explored the community structures inside different social environments, and found these community structures match quite well with the real underlying social structures.

Popularity Ranking. We observed the heterogeneity for node centralities in both global and local scales. We shall see that popular hubs are as useful in the PSN context as they are in the wireline Internet and in the web.

From Section 6 to Section 9, we will look at how we can use this information to make smart forwarding decisions. The following three preexisting schemes provide lower and upper bounds in terms of cost and delivery success. All of these schemes are inefficient because they assume a homogeneous environment. If the environment is homogeneous then every node is *statistically equivalent* (i.e., every node has the same likelihood of delivering the messages to the destination). As we showed in the first half of this paper, the environments and nodes are diverse and, hence, all these naive schemes are doomed to have poor performance. We need to design algorithms which make use of this rich heterogeneity.

- **WAIT:** Hold onto a message until the sender encounters the recipient directly, which represents the lower bound for delivery cost. WAIT is the only single-copy algorithm in this paper.
- **FLOOD:** Messages are flooded throughout the entire system, which represents the upper bound for delivery and cost.
- **Multiple-Copy-multiple-hoP (MCP):** Multiple copies are sent subject to a time-to-live hop count limit on the propagation of messages. By exhaustive emulations, the 4-copy-4-hop MCP scheme is found to be the most cost-effective scheme in terms of delivery ratio and cost for all naive schemes among most of the data sets. Hence, for fair comparison, we would

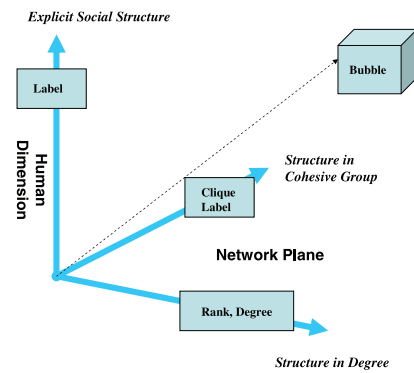


Fig. 5. Design space for forwarding algorithms.

like to evaluate our algorithms and the comparison algorithms (e.g., PROPHET and SimBet) against the 4-copy-4-hop MCP scheme in most of the cases.

Fig. 5 shows the design space for the forwarding algorithms. The vertical axis represents the explicit social structure. This is the social or human dimension. The two horizontal axes represent the network structural plane, which can be inferred purely from observed contact patterns. The Structure-in-Cohesive Group axis indicates the use of localized cohesive structure, and the Structure-in-Degree axis indicates the use of node ranking and degree. These are observable physical characteristics. In our design framework, it is not necessary that physical dimensions are orthogonal to the social dimension, but since they represent two different design parameters, we would like to separate them. The design philosophy here is to consider both the social and physical aspects of mobility.

We introduce four forwarding algorithms in this paper, namely LABEL, RANK, DEGREE, and BUBBLE.

- **LABEL:** Explicit labels are used to identify forwarding nodes that belong to the same organization. Optimizations are examined by comparing label of the potential relay nodes and the label of the destination node. This is in the human dimension, although an analogous version can be done by labeling a k -clique community in the physical domain.
- **RANK:** The forwarding metric used in this algorithm is the node centrality. A message is forwarded to nodes with higher centrality values than the current node. It is based on observations in the network plane, although it also reflects the hub popularity in the human dimension.
- **DEGREE:** The forwarding metric used in this algorithm is the node degree, more specifically the observed average of the degree of a node over a certain time interval. Either the last interval window (S-Window), or a long-term cumulative estimate, (C-Window) is used to provide a fully decentralized approximation for each node's centrality, and then that is used to select forwarding nodes.
- **BUBBLE:** The BUBBLE family of protocols combines the observed hierarchy of centrality of nodes and observed community structure with explicit labels, to decide on the best forwarding nodes. BUBBLE is an example algorithm that uses information from

both human aspects and also the physically observable aspects of mobility.

BUBBLE is a combination of LABEL and RANK. It uses RANK to spread out the messages and uses LABEL to identify the destination community. For this algorithm, we make two assumptions:

- Each node belongs to at least one community. Here, we allow single node communities to exist.
- Each node has a global ranking (i.e., global centrality) in the whole system and also a local ranking within its community. It may belong to multiple communities and, hence, may have multiple local rankings.

5.1 Hagglesim Emulator

In order to evaluate different forwarding algorithms, we developed an emulator called *Hagglesim* [20], which can replay the mobility traces collected and emulate different forwarding strategies on every contact event. This emulator is driven by contact events. The original trace files are divided into discrete sequential contact events, and fed into the emulator as inputs.

In all the simulations in this work, we divided the traces into discrete contact events with granularity of 100 s. Our simulator reads the file line by line, treating each line as a discrete encounter event, and makes a forwarding decision on this encounter based on the forwarding algorithm under study.

5.2 Simulation Parameters

There are three parameters we used in our simulation to achieve controlled flooding in MCP strategy.

- *Number of Copies (m)*. The maximum number of duplicates of each message created at each node.
- *Number of Hops (Hop-TTL)*. The maximum number of hops, counted from the source, that a message copy can travel before reaching the destination; this is similar to TTL value in the Internet.
- *Time TTL*. The maximum time a message can stay in the system after its creation. This is to prevent expired messages from further circulation.

For all the emulations conducted to compare forwarding efficiency in this paper, we have the following two metrics.

- *Delivery Ratio*. The proportion of messages that have been delivered out of the total unique messages created.
- *Delivery Cost*. The total number of messages (include duplicates) transmitted across the air. To normalize this, we divide it by the total number of unique messages created.

For some cases, we also compute the hop-count distribution for the deliveries, which is the distribution of the number of hops needed for all the deliveries, and which reveals the social distance between sources and destinations.

In the following sections, we will introduce several forwarding algorithms and evaluate their performances with the above metrics. For each emulation, 1,000 unique messages are created, with the source and destination randomly chosen or chosen based on specific grouping,




Fig. 6. Comparison of delivery ratio (left) and cost (right) of MCP and RANK on 4-copies and 4-hops case (*Reality*).

which will be specified in each evaluation case. The creating time of the messages is uniformly distributed within the simulation duration. Since the experimental durations for the data sets we used for simulation are in the order of days or even weeks, so we ignore the transient period and average the results throughout the simulation periods.

6 GREEDY RANKING ALGORITHM

The contribution of this section is to introduce RANK algorithm in detail and evaluate its performance using different data sets.

RANK is similar to the greedy strategy introduced by Adamic et al. [21]. A PSN is not a static network like the Internet; we do not know when a local maximum is reached since the next encounter is unexpected. We cannot employ precisely the same strategy as they proposed, of traversing up the hierarchy until reaching the maximum, and then down a step. In RANK, we assume each node knows only its own ranking and the rankings of those it encounters, but does not know the ranking of other nodes it does not encounter, and does not know which node has the highest rank in the system. RANK is very simple: We keep pushing traffic on all paths to nodes which have a higher ranking than the current node, until either the destination is reached, or the messages expire.

If a system is small enough, the global ranking of each node is actually the local ranking. If we consider only the Systems Research Group (around 40 people), a subset of the Cambridge Computer Laboratory (235 people), this is the ranking of each node inside the group. If we consider the whole Computer Laboratory, we are considering a larger system of many groups, but they all use the same building. A homogeneous ranking can still work. But when we consider the whole city of Cambridge, a homogeneous ranking system would exclude many small scale structures. In this section, we show that in relatively small and homogeneous systems, a simple greedy ranking algorithm can achieve good performance.

Fig. 6 (left) shows that RANK performs almost as well as MCP for delivery. Fig. 6 (right) also shows that the cost is at maximum of around 40 percent that of MCP, which represents a marked improvement. The audiences may notice that the difference in cost between these two algorithms is not constant. This is because they have different spreading mechanisms which may affect their abilities to find necessary number of relays within a certain time TTL.

RANK appears to work in small and homogeneous systems, but when we look at a more diversified system, for




Fig. 7. The hop distribution of the delivered (left) and the rank distribution of undelivered (right) on HK data.

example the *Hong Kong* data set, it may work differently. In the *Hong Kong* experiment, the 37 participants are intentionally selected without any social correlation. They live and work throughout the whole city. With FLOOD, we can deliver more than 40 percent of the total traffic across the whole city by using only the 37 iMotes and the external devices detected by these iMotes. But in this case, greedy ranking can only deliver 10 percent of the messages, although the cost is much lower as well. In terms of delivery and cost, greedy ranking is still more cost-effective than flooding, but clearly the delivery success rate is still too low. One explanation for this low performance is that since the participants have no social correlation, and belong to different social communities, high global ranking of a node may not represent a good choice of relay for some local communities. Messages keep being pushed up to some globally higher ranking nodes, and getting stuck at some maxima, rather than then trickling down to some local community.

Fig. 7a (left) shows that the maximum number of hops for greedy Rank is 4-hops and after that the messages get stuck. Fig. 7b (right) shows the rank distribution of the sources, destinations and dead ends of all the undelivered messages, indicating that message delivery has typically failed at highly-ranked nodes. This supports our hypothesis concerning the dilemma of the messages getting stuck at maxima.

Hierarchical organization is a common feature of many complex systems. The defining feature of a hierarchical organization is the existence of a hierarchical path connecting any two of its nodes. Trusina et al. [22] address how to detect and measure the extent of the hierarchy manifested in the topology of a given complex network. They defined the hierarchical path based on node degrees. A path between two nodes in a network is called hierarchical if it consists of an *up path*, where one is allowed to step from node i to node j only if their degrees k_i, k_j satisfy $k_i \leq k_j$, followed by a *down path* where only steps to nodes of lower or equal degree are allowed. Either the up or down path is allowed to have zero length. Because of the good results from the greedy ranking algorithm, we analyzed the percentage of hierarchical paths inside all the shortest paths. Table 4 summarizes the results.

The percentage of hierarchical paths is calculated as the number of hierarchical paths divided by the number of nondirect deliveries. We can see that for *Cambridge* data and *Reality*, the percentage of hierarchical paths is very high, so our strategy of pushing the messages up the ranking tree can find a lot of these paths, and the performance of the

TABLE 4
Hierarchical Paths Analysis of All Shortest Paths

Experimental dataset	% hierarchical paths
Cambridge	87.2 (-2.4,+4.3)
Reality	81.9 (-3.1,+3.3)
Infocom05	62.3 (-2.5,+2.5)
Infocom06	69.5 (-4.1,+2.4)
Hong Kong	33.5 (-4.0,+4.0)

ranking strategy here is not much different from that of MCP. For *Infocom06* and *Infocom05*, the percentage of hierarchical paths is also high, hence the greedy RANK strategy can as well discover many of the shortest paths. However, for the *Hong Kong* experiment, the network is too sparse and a lot of shortest paths are hidden. (This occurs because we could not know the devices detected by the external devices, and most of the resulting paths used for delivery are actually not the shortest.) We can see that percentage of hierarchical paths controls the delivery success achieved by the greedy RANK algorithm. We conclude from this that a high percentage of the shortest paths are actually hierarchical paths.

7 DIRECT LABELLING STRATEGY

In the LABEL strategy, each node is assumed to have a label that tells others its affiliation, just like a name badge in a conference. The direct LABEL strategy refers to the exclusively using of labels to forward messages to destinations: Next-hop nodes are selected if they belong to the same group (same label) as the destination. It was demonstrated that LABEL significantly improves forwarding efficiency over “oblivious” forwarding using *Infocom06* data set [20]. This is a beginning of social-based forwarding in PSN. The limitation of LABEL is the lack of mechanisms to move messages away from the source when the destinations are socially far away (such as *Reality*). The contribution of this section is to demonstrate the limitations of LABEL strategy, which motivates a new forwarding algorithm using both community and centrality information.

We evaluate the LABEL strategy on the *Reality* data set. Here, we use the community information detected using *K-CLIQUE* algorithm to label the nodes. We can see from Fig. 8 that LABEL only achieves around 55 percent of the delivery ratio of the MCP strategy and only 45 percent of the flooding delivery although the cost is also much lower.




Fig. 8. Comparison of delivery ratio (left) and cost (right) of MCP and LABEL on 4-copies and 4-hops case (*Reality*).




Fig. 9. Correlation of n th-hop nodes with the source group and destination group (*Reality*).

However, it is not an ideal scenario for LABEL. In this environment, people do not mix as well as in a conference. A person in one group may not meet members in another group so often, so waiting until the members of the other group appear to do the transmission is not effective here.

Fig. 9 shows the correlation of the n th-hop relay nodes to the source and destination groups for the messages on all the shortest paths, that is the percentage of the n th-hop relay nodes that are still in the same group as the source or already in the same group as the destination. We can see that more than 50 percent of the nodes on the first hops (from the S-Group plot) are still in the same group as the source group of the message and only around 5 percent of the first hop nodes (from the D-Group plot) are in the same group as the destination. We can also see that on going to the 2nd hop, S-Group correlation drops to slightly less than 30 percent, and when going to 4th-hops, almost all (90 percent) messages have escaped from this source group. To calculate the percentage for each hop, we divide the number of messages which belong to that group (S-Group or D-Group) by the total number of messages destined beyond nodes at that particular hop, but not the total messages created. In the 4-hop case, there are perhaps only 100 messages to forward further, and only 10 out of these 100 relay nodes belong to the source group. This explains why LABEL is not effective, since it is far from discovering the shortest path. In the next section, we will talk about how to use centrality to improve the delivery ratio of LABEL.

8 CENTRALITY MEETS COMMUNITY

The contribution of this section is to combine the knowledge of both centralities of nodes and community structure, to achieve further performance improvements in forwarding. We show that this avoids the occurrence of the dead ends encountered with pure global ranking schemes. We call the protocols here BUBBLE, to capture our intuition about the social structure. Messages bubble up and down the social hierarchy, based on the observed community structure and node centrality, together with explicit label data. Bubbles represent a hybrid of social and physically observable heterogeneity of mobility over time and over community.

8.1 Two-Community Case

In order to make the study more systematic, we start with the two-community case. We use the *Cambridge* data set for




Fig. 10. Node centrality in 2 groups in *Cambridge* data (see Section 4 for the method of calculating the centrality values). (a) Group A. (b) Group B.

this study. By experimental design, and confirmed using our community detection algorithm, we can clearly divide the *Cambridge* data into two communities—the undergraduate year-one and year-two group. In order to make the experiment more fair, we limit ourselves to just the two 10-clique groups found with a number-of-contact threshold of nine, that is, where each node at least meet another nine nodes frequently. Some students may skip lectures and cause variations in the results, so this limitation makes our analysis yet more plausible.

First we look at the simplest case, for the centrality of nodes within each group. In this case, the traffic is created only for members within the same community and only members in the same community are chosen as relays for messages. We can clearly see from Figs. 10a and 10b that inside a community, the centrality of each node is different. In Group B, there are two nodes which are very popular, and have relayed most of the traffic. All the other nodes have low centrality value. Forwarding messages to the popular nodes would make delivery more cost effective for messages within the same community.

Then we consider traffic which is created within each group and only destined for members in another group. To eliminate other outside factors, we use only members from these two groups as relays. Fig. 11 shows the individual node centrality when traffic is created from one group to another and the correlation of node centrality within an individual group and intergroup (for data deliveries only to other groups but not to its only group) centrality. We can see that points lie more or less around the diagonal line. This means that the inter- and intragroup centralities are quite well correlated. Active nodes in a group are also active nodes for intergroup communication. There are some




Fig. 11. Intergroup centrality (left) and correlation between intra- and intergroup centrality (right), *Cambridge*.




Fig. 12. Correlation of local centrality of Group A and the global centrality (Cambridge).

points on the left hand side of the graph which have low intragroup centrality but moderate intergroup centrality. These are nodes which move across groups. They are not important for intragroup communication but can perform certainly well when we need to move traffic from one group to another.

We can show now why homogeneous global ranking in Section 6 does not work perfectly. Fig. 12 shows the correlation of the local centrality of Group A and the global centrality of the whole population. We can see that quite a number of nodes from Group A lie along the diagonal line. In this case, the global ranking can help to push the traffic toward Group A. However, the problem is that some nodes which have very high global rankings are actually not members of Group A, for example node D. Just as in real society, a politician could be very popular in the city of Cambridge, but not a member of the Computer Laboratory, so may not be a very good relay to deliver message to the member in the Computer Laboratory. Now we assume there is a message at node A to deliver to another member of Group A. According to global ranking, we would tend to push the traffic toward B, C, D, and E in the graph. If we pushed the traffic to node C, it would be fine, and to node B it would be perfect. But if it pushed the traffic to node D and E, the traffic could get stuck there and not be routed back to Group A. If it reaches node B, that is the best relay for traffic within the group, but node D has a higher global ranking than B, and would tend to forward the traffic to node D, where it would probably get stuck again. Here, we propose the BUBBLE algorithm to avoid these dead ends.

Forwarding is carried out as follows: If a node has a message destined for another node, this node would first bubble this message up the hierarchical ranking tree using the global ranking until it reaches a node which has the same label (community) as the destination of this message. Then the local ranking system will be used instead of the global ranking and continue to bubble up the message through the local ranking tree until the destination is reached or the message expired. This method does not require every node to know the ranking of all other nodes in the system, but just to be able to compare ranking with the node encountered, and to push the message using a greedy approach. We call this algorithm BUBBLE, since each world/community is like a bubble. Fig. 13 illustrates the




Fig. 13. Illustration of the BUBBLE forwarding algorithm.

BUBBLE algorithm and the pseudocode can be found in our previous work [18].

This fits our intuition in terms of real life. First you try to forward the data via people more popular than you around you, and then bubble it up to well-known popular people in the society, such as a postman. When the postman meets a member of the destination community, the message will be passed to that community. This community member will try to identify the more popular members within the community and bubble the message up again within the local hierarchy until the message reaching a very popular member, or the destination itself, or the message expires.

A modified version of this strategy is that whenever a message is delivered to the community, the original carrier can delete this message from its buffer to prevent it from further dissemination. This assumes that the community member would be able to deliver this message. We call this protocol with deletion, strategy BUBBLE-B, and the original algorithm introduced above BUBBLE-A.

We can see from Fig. 14 that both BUBBLE-A and BUBBLE-B achieve almost the same delivery success rate as the 4-copy-4-hop MCP. Although BUBBLE-B has the message deletion mechanism, it achieves exactly the same delivery as BUBBLE-A. BUBBLE-A only has 60 percent the cost of MCP and BUBBLE-B is even better, with only 45 percent the cost of MCP. Both have almost the same delivery success as MCP.

8.2 Multiple-Community Cases

To study the multiple-community cases, we use the *Reality* data set. To evaluate the forwarding algorithm, we extract a three-week session during term time from the whole

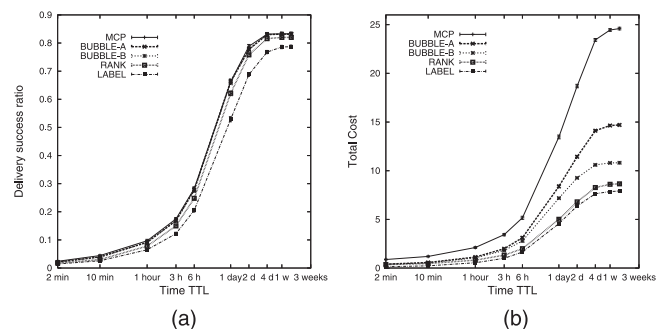


Fig. 14. Comparisons of several algorithms on Cambridge data set.




Fig. 15. Node centrality in several individual groups (*Reality*).




Fig. 16. Comparisons of several algorithms on *Reality* data set, single group.

nine-month data set. Emulations are run over this data set with uniformly generated traffic.

There is a total of eight groups within the whole data set. Fig. 15 shows the node centrality in four groups, from small-size to medium-size and large-size group. We can see that within each group, almost every node has different centrality.

In order to make our study easier, we first isolate the largest group in Fig. 15, consisting of 16 nodes. In this case, all the nodes in the system create traffic for members of this group. We can see from Fig. 16 that BUBBLE-A and BUBBLE-B perform very similarly to MCP most of the time in the single group case, and even outperform MCP when the time TTL is set to be larger than one week. BUBBLE-A only has 70 percent and BUBBLE-B only 55 percent of the cost of MCP. We can say that the BUBBLE algorithms are much more cost effective than MCP, with high delivery ratio and low delivery cost.

After the single group case, we start looking at the case of every group creating traffic for other groups, but not for its own members. We want to find the upper cost bound for the BUBBLE algorithm, so we do not consider local ranking (i.e., only global ranking); messages can now be sent to all members in the group. This is exactly a combination of direct LABEL and greedy RANK, using greedy RANK to move the messages away from the source group. We do not implement the mechanism to remove the original message after it has been delivered to the group member, so the cost here will represent an upper bound for the BUBBLE algorithms.

From Fig. 17, we can see that of course flooding achieves the best performance for delivery ratio, but the cost is 2.5 times that of MCP, and 5 times that of BUBBLE. BUBBLE




Fig. 17. Comparisons of several algorithms on *Reality* data set, all groups.




Fig. 18. Comparisons of several algorithms on *Infocom06* data set, all groups.

is very close in performance to MCP in multiple groups case as well, and even outperforms it when the time TTL of the messages is allowed to be larger than two weeks.² However, the cost is only 50 percent that of MCP. Fig. 18 shows the same performance evaluations with the *Infocom06* data set. In this case, the delivery ratio of RANK is approaching that of MCP but with less than half of the cost. The performance of BUBBLE over RANK is not as significant as in the *Reality* case because in a conference scenario the people are very mixing and, hence, the community factors are less dominating. We can also see that even in this case, the delivery cost for BUBBLE only increases slightly, which indicates that even in a mixing environment, BUBBLE is still very robust toward the possible misleading of the community factors.

In order to further justify the significance of social-based forwarding, we also compare BUBBLE with a benchmark “nonoblivious” forwarding algorithm, PROPHET [5], and another state-of-the-arts social-based forwarding algorithm, SimBet [10]. PROPHET uses the history of encounters and transitivity to calculate the probability that a node can deliver a message to a particular destination. SimBet is similar in concept as BUBBLE for leveraging social contexts. It exploits the exchange of preestimate “betweenness” centrality metrics and locally determined social “similarity” to the destination node to guide the message delivery. Since PROPHET has been evaluated against other algorithm before and SimBet is another well-credited social-based algorithm, and both have the same contact-based nature as BUBBLE (i.e., does not need location information), they are good candidates to compare with BUBBLE.

2. Two weeks seems to be very long, but as we have mentioned before, the *Reality* network is very sparse. We choose it mainly because it has long experimental period and hence more reliable community structures can be inferred. The evaluations here can serve as a proof of concept of the BUBBLE algorithm, although the delays are large in this data set.




Fig. 19. Comparisons of BUBBLE, PROPHET, and SimBet on *Reality* data set.

PROPHET has four parameters. We use the default PROPHET parameters as recommended in [5]. However, one parameter that should be noted is the time elapsed unit used to age the contact probabilities. The appropriate time unit used differs depending on the application and the expected delays in the network. Here, we age the contact probabilities at every new contact. In a real application, this would be a more practical approach since we do not want to continuously run a thread to monitor each node entry in the table and age them separately at different time. For SimBet routing, we use the default values for the SimBet utility parameters as specified in the paper [10] (i.e., $\alpha = \beta = 0.5$), which assigns an equal importance to the similarity and betweenness utility. And since BUBBLE is a multicopy algorithm, we also implement a multicopy SimBet for better comparison. For all the comparisons, we use the same settings of number of hops and number of copies for all the algorithms, more particularly we show the results of the 4-hop-4-copy case. The results are also valid for the other combinations of number of hops and number of copies in our simulations. As most of the traces have long experimental durations (i.e., in the order of days or even weeks), we average the results throughout the simulation period and ignore the transient period.

Fig. 19 shows the comparison of the delivery ratio and delivery cost of BUBBLE, PROPHET, and SimBet for the 4-hop-4-copy case.³ Here, for the delivery cost, we only count the number of copies created in the system for each message, as we have done before for the comparison with the “oblivious” algorithms. We do not count the control traffic created by PROPHET for exchanging routing table during each encounter, which can be huge if the system is large (PROPHET uses flat addressing for each node and its routing table contains entry for each known node). We also do not count the message exchange in SimBet for updating the similarity and betweenness values. We can see that most of the time, BUBBLE achieves a similar delivery ratio to PROPHET and around 10 percent better than SimBet, but with only half of the cost of PROPHET and 70 percent of the cost of SimBet. Considering that BUBBLE does not need to keep and update an routing table for each node pairs, the performance achievement is significant.

PROPHET relies on encountering history and transient delivery predictability to choose relays. This can efficiently identify the routing paths to the destinations, but the dynamic environment may result in many nodes having a

3. For one more perspective on the data, here we show the 4-hop-4-copy case instead of the unlimited case in Fig. 17. The trend is very similar for the other case.




Fig. 20. Correlation of rank with total degree and rank with unit time degree (*Reality*).

lot of slightly fluctuation of probabilities. This results in more redundant nodes being chosen as relays, which can be reflected from the delivery cost. Instead, BUBBLE uses social information and, hence, filters out these noises due to the temporal fluctuations of the network. SimBet can successfully leverage social context, but it fails in identifying the sequence of using betweenness and similarity. BUBBLE explicitly identifies centrality and community, and first uses centrality metric to spread out the messages and then uses community metric to focus the messages to the destinations. This approach effectively guarantees a high delivery ratio and a low delivery cost.

A remark here is that the centrality values used for the BUBBLE simulations in this section are calculated in a centralized way, while PROPHET and SimBet use mainly online estimation, but we will show that this can be effectively approximated in a low-cost distributed manner in the next section. Overall, we evaluated BUBBLE against WAIT, FLOOD, the optimized MCP, LABEL, RANK, the benchmark PROPHET, and SimBet. This provides us a reasonable variety of samples to illustrate the performance of BUBBLE.

9 MAKING CENTRALITY PRACTICAL

For practical applications, we want to look further into how BUBBLE can be implemented in a distributed way. To achieve this, each device should be able to detect its own community and calculate its centrality values. In [23], we have proposed three algorithms, named SIMPLE, K -CLIQUE, and MODULARITY for distributed community detection, and we have proved that detection accuracy can be up to 85 percent of the centralized K -CLIQUE algorithm. The next step is to ask how each node can know its own centrality in a decentralized way, and how well past centrality can predict the future.

The final contribution of this paper is to provide answers to these two questions.

9.1 Approximating Centrality

We found that the total degree (unique nodes) seen by a node throughout the experiment period is not a good approximation for node centrality. Instead, the degree per unit time (for example the number of unique nodes seen per 6 hours, started from midnight) and the node centrality have a high correlation value. We can see from Fig. 20 that some nodes with a very high total degree are still not good carriers. It also shows that the per 6 hour degree is quite well correlated to the centrality value, with correlation coefficient as high as 0.9511. That means how many people




Fig. 21. Comparisons of delivery (left) and cost (right) of RANK, S-Window and C-Window (*Reality*).

you know does not matter too much, but how frequently you interact with these people does matter.

In order to verify that the average unit-time degree is as good as or close to RANK, we run another set of emulations using greedy average unit-time degree (or we simply call it DEGREE) instead of the precalculated centrality. We find that RANK and DEGREE perform almost the same with the delivery and cost lines overlapping each other. They not only have similar delivery but also similar cost.

However, the average unit-time degree calculated throughout the whole experimental period is still difficult for each node to calculate individually. We then consider the degree for the previous unit-time slot (we call this the slot window) such that when two nodes meet each other, they compare how many unique nodes they have met in the previous unit-time slot (e.g., 6 hours). We call this approach single window (S-Window). Another approach is to calculate the average value on all previous windows, such as from yesterday to now, then calculate the average degree for every 6 hours. We call this approach cumulative window (C-Window). This technique is similar to a statistics technique called exponential smoothing [24] and we would like to do further theoretical investigation.

We can see from Fig. 21 that the S-Window approach reflects more recent context, and achieves a maximum of 4 percent improvement in delivery ratio over RANK, but at double the cost. The C-Window approach measures more of the cumulative effect, and gives more stable statistics about the average activeness of a node. However, its cumulative measurement is not as good an estimate as RANK, which averages throughout the whole experimental period. It does not achieve as good delivery as RANK (not more than 10 percent less in term of delivery), but it also has lower cost.

9.2 Predictability of Centrality

In order to further verify whether the centrality measured in the past is useful as a predictor for the future, we extracted three temporally consecutive three-week sessions from the *Reality* data set and then run a set of greedy RANK emulations on the last two data sessions, but using the centrality values from first session.

Fig. 22 shows the delivery ratio and cost of RANK on the second data session using the centrality values from the first data session. It seems that the performance of RANK is not far from MCP but with much lower cost, i.e., it is as good as running the emulation on the original data set, which the centrality values derived from. Similar performance is also observed in the third data session. These results imply some




Fig. 22. Delivery ratio (left) and cost (right) of RANK algorithm on second data session, all groups (*Reality*).

level of predictability of human mobility, and show empirically that past contact information can be used in the future.

All these approaches, (DEGREE, S-Window, C-Window and predictability of human mobility) provide us with a decentralized way to approximate the centrality of nodes in the system, and hence help us to design appropriate forwarding algorithms. Combining these approximate methods and the distributed community detection, we can put BUBBLE into reality. We will briefly discuss how distributed BUBBLE works for a citywide environment, but leave the evaluation details as future work when we can get a larger scale of data set.

Suppose there is a network of mobile users, perhaps spanning an entire city, each device can detect its own local community or knowing its social graph from online social networks [25]. At the same time, it also counts its own 6-hour-averaged degree (i.e., C-Window). Its global ranking can be approximated as its 6-hour-averaged degree for all nodes and its local ranking can be approximated as its 6-hour-averaged degree only for nodes inside its community. With all these metrics, each node can forward messages using BUBBLE. Or, we simply call it DiBuBB algorithm, which uses labels from its social graph, affiliation, or distributed community detection for community information and C-Window to approximate its own global and local centrality values. Besides that, it operates exactly like BUBBLE. Fig. 23 shows the plotting of delivery ratio against the delivery cost for BUBBLE, DiBuBB,




Fig. 23. Delivery ratio against cost for several algorithms (*Reality*).

SimBet, and PROPHET. Here, DiBuBB uses the 6-hour C-Window approach to approximate the centrality values and the social graph information for the communities. In general, the larger the slopes of the lines, the more efficient the algorithm is, in term of delivery and cost. We can see that the performance of DiBuBB is very close to BUBBLE and outperforms both SimBet, and PROPHET.

10 RELATED WORK

For distributed search for nodes and content in power-law networks, Sarshar et al. [26] proposed using a probabilistic broadcast approach: sending out a query message to an edge with probability just above the bond⁴ percolation threshold of the network. They show that if each node caches its directory via a short random walk, then the total number of accessible contents exhibits a first-order phase transition, ensuring very high hit rates just above the percolation threshold.

For routing and forwarding in DTNs and mobile ad hoc networks, there is much existing literature. Vahdat et al. proposed epidemic routing, which is similar to the “oblivious” flooding scheme we evaluated in this paper [27]. Spray and Wait is another “oblivious” flooding scheme but with a self-limited number of copies [28]. Grossglauser et al. proposed the two-hop relay schemes to improve the capacity of dense ad hoc networks [29]. Many approaches calculate the probability of delivery to the destination node, where the metrics are derived from the history of node contacts, spatial information and so forth. The pattern-based Mobospace Routing by Leguay et al. [30], location-based routing by Lebrun et al. [31], context-based forwarding by Musolesi et al. [32] and PROPHET Routing [5] fall into this category. PROPHET uses past encounters to predict the probability of future encounters. The transitive nature of encounters is exploited, where indirectly encountering the destination node is evaluated. Message Ferry by Zhao et al. [33] takes a different approach by controlling the movement of each node.

Recent attempts to uncover a hidden stable network structure in DTNs such as social networks have been emerged. For example, SimBet Routing [10] uses egocentric centrality and its social similarity. Messages are forwarded toward the node with higher centrality to increase the possibility of finding the potential carrier to the final destination. LABEL forwarding [20] uses affiliation information to help forwarding in PSNs based on the simple intuition that people belonging to the same community are likely to meet frequently, and thus act as suitable forwarders for messages destined for members of the same community. We have compared BUBBLE with LABEL and SimBet in this paper, and demonstrated that by the exploitation of both community and centrality information, BUBBLE provided further improvement in forwarding efficiency. The mobility-assisted Island Hopping forwarding [34] uses network partitions that arise due to the distribution of nodes in space. Their clustering approach is based on the significant locations for the nodes and not for clustering nodes

themselves. Clustering nodes is a complex task to understand the network structure for aid of forwarding.

Finally, we emphasize that we take an experimental rather than theoretical approach, which contrasts with other work described above.

11 CONCLUSION

We have shown that it is possible to uncover important characteristic properties of social network from a diverse set of real world human contact traces. We have demonstrated that community and centrality social metrics can be effectively used in forwarding decisions. Our BUBBLE algorithm is designed for a delay-tolerant network environment, built out of human-carried devices, and we have shown that it has similar delivery ratio to, but much lower resource utilization than flooding, control flooding, PROPHET, and SimBet.

BUBBLE is designed to work better with a hierarchical community structure. The limitation imposed by the size of the data sets (each experiment is not large enough for us to extract hierarchical structure) does not allow us to optimally evaluate it. The current evaluation on a flat community structure did still provide us satisfactory performance improvement. We will further verify our results when more mobility traces are available. Synthetic mobility models can also be useful for further evaluating the algorithm, but currently there is no benchmark models. Another aspect we want to look using our mobility traces is to compare them with the available mobility models and find out the one which can represent most of real mobility scenarios. We believe that this approach represents an early step in combining rich multilevel information of social structures and interactions to drive novel and effective means for disseminating data in DTNs. A great deal of future research can follow.

ACKNOWLEDGMENTS

This research was funded in part by the Deutsche Telekom Laboratories MADNet project, the ITA project, the EU grants for the Hagggle project, IST-4-027918, and the SOCIALNETS project, 217141. The authors acknowledge comments from Steven Hand, Brad Karp, Frank Kelly, Richard Mortier, Pietro Lio, Andrew Moore, Nishanth Sastry, Derek Murray, Sid Chau, Andrea Passarella, Hamed Haddadi, and Georgios Smaragdakis, and SimBet source codes from Thrasyvoulos Spyropoulos.

REFERENCES

- [1] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, “Pocket Switched Networks and Human Mobility in Conf. Environments,” *Proc. ACM Special Interest Group Data Comm. Workshop (SIGCOMM '05)*, 2005.
- [2] K. Fall, “A Delay-Tolerant Network Architecture for Challenged Internets,” *Proc. ACM Special Interest Group Data Comm. Workshop (SIGCOMM '03)*, 2003.
- [3] D. Kempe, J. Kleinberg, and A. Kumar, “Connectivity and Inference Problems for Temporal Networks,” *J. Computer and System Sciences*, vol. 64, no. 4, pp. 820-842, 2002.
- [4] E.P.C. Jones, L. Li, and P.A.S. Ward, “Practical Routing in Delay-Tolerant Networks,” *Proc. ACM Special Interest Group Data Comm. Workshop (SIGCOMM '05)*, 2005.

4. A percolation which considers the lattice edges as the relevant entities.

- [5] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic Routing in Intermittently Connected Networks," *ACM SIGMOBILE Mobile Computing and Comm. Rev.*, vol. 7, no. 3, July 2004.
- [6] S. Okasha, "Altruism, Group Selection and Correlated Interaction," *British J. for the Philosophy of Science*, vol. 56, no. 4, pp. 703-725, Dec. 2005.
- [7] M.E.J. Newman, "Detecting Community Structure in Networks," *The European Physical J. B*, vol. 38, pp. 321-330, 2004.
- [8] L. Danon, J. Duch, A. Diaz-Guilera, and A. Arenas, "Comparing Community Structure Identification," *J. Statistical Mechanics*, vol. 2005, no. 9, p. P09008, Sept. 2005.
- [9] L.C. Freeman, "A Set of Measuring Centrality Based on Betweenness," *Sociometry*, vol. 40, pp. 35-41, 1977.
- [10] E. Daly and M. Haahr, "Social Network Analysis for Routing in Disconnected Delay-Tolerant Manets," *Proc. ACM MobiHoc*, 2007.
- [11] N. Eagle and A. Pentland, "Reality Mining: Sensing Complex Social Systems," *Personal and Ubiquitous Computing*, vol. 10, no. 4, pp. 255-268, May 2006.
- [12] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms," *Proc. IEEE INFOCOM*, Apr. 2006.
- [13] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnović, "Power Law and Exponential Decay of Inter Contact Times between Mobile Devices," *Proc. ACM MobiCom*, 2007.
- [14] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft, "Opportunistic Content Distribution in an Urban Setting," *Proc. ACM MobiCom Workshop Challenged Networks (CHANTS '06)*, pp. 205-212, 2006.
- [15] G. Palla, I. Derényi, I. Farkas, and T. Vicsek, "Uncovering the Overlapping Community Structure of Complex Networks in Nature and Society," *Nature*, <http://dx.doi.org/10.1038/nature03607>, vol. 435, no. 7043, pp. 814-818, 2005.
- [16] M.E.J. Newman, "Analysis of Weighted Networks," *Physical Rev. E*, vol. 70, p. 056131, 2004.
- [17] P. Hui and J. Crowcroft, "Human Mobility Models and Opportunistic Communications System Design," *Philosophical Trans. Royal Soc. A: Math., Physical and Eng. Sciences*, vol. 366, no. 1872, pp. 2005-2016, June 2008.
- [18] P. Hui, J. Crowcroft, and E. Yoneki, "Bubble Rap: Social-Based Forwarding in Delay Tolerant Networks," *Proc. ACM MobiHoc*, May 2008.
- [19] M.E.J. Newman and M. Girvan, "Finding and Evaluating Community Structure in Networks," *Physical Rev. E*, <http://arxiv.org/abs/cond-mat/0308217>, vol. 69, Feb. 2004.
- [20] P. Hui and J. Crowcroft, "How Small Labels Create Big Improvements," *Proc. IEEE Int'l Workshop Intermittently Connected Mobile Ad Hoc Networks (ICMAN '07)*, Mar. 2007.
- [21] L.A. Adamic, B.A. Huberman, R.M. Lukose, and A.R. Puniyani, "Search in Power Law Networks," *Physical Rev. E*, vol. 64, pp. 46 135-46 143, Oct. 2001.
- [22] A. Trusina, S. Maslov, P. Minnhagen, and K. Sneppen, "Hierarchy Measures in Complex Networks," *Physical Rev. Letters*, <http://dx.doi.org/10.1103/PhysRevLett.92.178702>, vol. 92, p. 178702, 2004.
- [23] P. Hui, E. Yoneki, S.-Y. Chan, and J. Crowcroft, "Distributed Community Detection in Delay Tolerant Networks," *Proc. Second ACM/IEEE Int'l Workshop Mobility in the Evolving Internet Architecture (MobiArch '07)*, Aug. 2007.
- [24] P. Winters, "Forecasting Sales by Exponentially Weighted Moving Averages," *Management Science*, vol. 6, pp. 324-342, 1960.
- [25] P. Hui and N. Sastry, "Real World Routing Using Virtual World Information," *Proc. Int'l Conf. Computational Science and Eng. (CCSE '09)*, pp. 1103-1108, 2009.
- [26] N. Sarshar, P.O. Boykin, and V.P. Roychoudhury, "Scalable Percolation Search in Power Law Networks: Making Unstructured Peer-to-Peer Networks Scalable," *Proc. Fourth Int'l Conf. Peer-to-Peer Computing (P2P '04)*, June 2004.
- [27] A. Vahdat and D. Becker, "Epidemic Routing for Partially Connected Ad Hoc Networks," Technical Report CS-200006, Duke Univ., Apr. 2000.
- [28] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks," *Proc. ACM Special Interest Group Data Comm. Workshop Delay-Tolerant Networking (WDTN '05)*, 2005.
- [29] M. Grossglauser and D. Tse, "Mobility Increases the Capacity of Ad-Hoc Wireless Networks," *IEEE/ACM Trans. Networks*, vol. 10, no. 4, pp. 477-486, Aug. 2002.
- [30] J. Leguay, T. Friedman, and V. Conan, "Evaluating Mobility Pattern Space Routing for DTNs," *Proc. IEEE INFOCOM*, 2006.
- [31] J. Lebrun, C.-N. Chuah, D. Ghosal, and M. Zhang, "Knowledge-based Opportunistic Forwarding in Vehicular Wireless Ad Hoc Networks," *Proc. Vehicular Technology Conf.*, vol. 4, pp. 2289-2293, 2005.
- [32] M. Musolesi, S. Hailes, and C. Mascolo, "Adaptive Routing for Intermittently Connected Mobile Ad Hoc Networks," *Proc. Sixth IEEE Int'l Symp. World of Wireless Mobile and Multimedia Networks (WOWMOM '05)*, pp. 183-189, 2005.
- [33] W. Zhao, M. Ammar, and E. Zegura, "A Message Ferrying Approach for Data Delivery in Sparse Mobile Ad Hoc Networks," *Proc. ACM MobiCom*, pp. 187-198, 2004.
- [34] M.P.N. Sarafijanovic-Djukic and M. Grossglauser, "Island Hopping: Efficient Mobility-Assisted Forwarding in Partitioned Networks," *Proc. 3rd Ann. IEEE Comm. Soc. Sensor and Ad Hoc Comm. Networks (SECON '06)*, vol. 1, pp. 226-235, 2006.



Pan Hui received the bachelor's and Mphil degrees from the University of Hong Kong and the PhD degree from the Computer Laboratory, University of Cambridge, in 2007. While pursuing the PhD degree, he was also affiliated with Intel Research Cambridge and Thomson Research Laboratory in Paris. Currently, he is a senior research scientist at Deutsche Telekom Laboratories, Berlin, Germany. His research interests include delay tolerant networking, mobile network and systems, planet-scale mobility measurement, social networks, and the application of complex network science in communication system design. He is a member of the IEEE. More information can be found at <http://www.deutsche-telekom-laboratories.de/panhui>.



Jon Crowcroft graduated in physics from Trinity College, Cambridge University, United Kingdom, in 1979, and received the MSc degree in computing in 1981 and the PhD degree in 1993 from University College London (UCL), United Kingdom. He is currently the Marconi Professor of Communications Systems in the Computer Lab at the University of Cambridge, United Kingdom. He is a fellow of the ACM, the British Computer Society, the IE[ET] the Royal Academy of Engineering, and the IEEE. Jon was a recipient of the ACM Sigcomm Award in 2009.



Eiko Yoneki received a postgraduate diploma in computer science and the PhD degree in 2002 and 2006, respectively, from the University of Cambridge, United Kingdom. She is currently an EPSRC research fellow in the University of Cambridge Computer Laboratory, Systems Research Group, United Kingdom. Previously, she spent several years with IBM (US, Japan, Italy, and UK) working on various networking products. Her research interests span distributed systems, networking, and databases over mobile/wireless networks including complex and time-dependent networks. She is a member of the IEEE. More information can be found at <http://www.cl.cam.ac.uk/ey204>.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.