

BUBBLE Rap: Social-based Forwarding in Delay Tolerant Networks

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

Abstract—In this paper we seek to improve our understanding of human mobility in terms of social structures, and use these structures in the design of forwarding algorithms for Pocket Switched Networks (PSNs). Taking human mobility traces from the real world, we discover that human interaction is heterogeneous both in terms of hubs (popular individuals) and groups or communities. We propose a social based forwarding algorithm, BUBBLE, which is shown empirically to improve the forwarding efficiency significantly compared to oblivious forwarding schemes and to the PROPHET algorithm. We also show how this algorithm can be implemented in a distributed way, which demonstrates that it is applicable to the decentralised systems such as PSNs.

Index Terms—Social Network, Forwarding, Delay Tolerant Network, Pocket Switched Network, Community, Centrality.

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 uses contact opportunities to allow humans to communicate without network infrastructure. We require an efficient data forwarding mechanism over the temporal graph of the PSN [3], that copes with dynamical, repeated disconnection and re-wiring. End-to-end delivery through traditional routing algorithms is rarely applicable.

Many MANET and some DTN routing algorithms [4] [5] provide forwarding by building and updating routing tables whenever mobility occurs. We believe this approach is not cost effective for a PSN, since mobility is often unpredictable, and topology changes can be rapid. Rather than exchange much control traffic to create unreliable routing structures, we prefer to search for some characteristics of the network which are less volatile than mobility. A PSN is formed by people. Those people's social relationships form social topology, which may vary but such topology must be stable, and therefore can be used for better forwarding decisions. Furthermore, if we can detect these social mobility patterns online in a decentralised way, we can apply these algorithms in practice.

In this paper, we focus on two specific aspects of society: community and centrality. *Community* is an important attribute of PSNs. Cooperation binds, but also divides human society into communities. Human society is structured. Within a community, some people are more popular, and interact with more people than others (i.e. have high *centrality*); we call them hubs. Popularity ranking is one aspect of the population. 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.

Methodologically, community detection [7] [8] can help us to 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 these questions:

- 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 decentralised way for such schemes to operate?

To preview our answers to the above questions, we evaluate the impact of community and centrality on forwarding, and propose a hybrid algorithm, BUBBLE, that selects high centrality nodes and community members of destination as relays. We demonstrate a significant improvement in forwarding efficiency over oblivious forwarding and the PROPHET algorithm [5], which uses patterns of movement, rather than the longer term social relationships on which our BUBBLE scheme rests. In a PSN, there may be no a priori information. By definition, we are also in a decentralised 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

- P. Hui is with Deutsche Telekom Laboratories, Ernst-Reuter-Platz 7, 10587, Berlin, Germany
E-mail: see <http://www.michaelshell.org/contact.html>
- J. Crowcroft and E. Yoneki are with Computer Laboratory, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge, CB3 0FD, UK

of packet 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, and then describe the contact graphs in Section 3 and inferring human communities in Section 4. In Section 5, we examine the heterogeneity in centrality. We show evaluation methodology in Section 6 and results including discussions from Section 7 to 11 followed by the related works in Section 12. Finally, we conclude the paper a brief summary of contributions in Section 13.

2 EXPERIMENTAL DATASETS

In this paper, we use three experimental datasets gathered by the Huggle Project¹ over two years, referred to as *HongKong*, *Cambridge*, *Infocom06*; one dataset from the MIT Reality Mining Project [10], referred to as *Reality*. Previously, the characteristics of these datasets such as inter-contact and contact distribution have been explored in several studies [11] [1] [12], to which we refer the reader for further background information. These four datasets cover a rich diversity of environments, ranging from busy metropolitan city (*HongKong*) to quite university town (*Cambridge*), with an experimental period from several days (*Infocom06*) to almost one year (*Reality*).

- In *Infocom05*, the devices were distributed to approximately fifty students attending the Infocom student workshop. Participants belong to different social communities (depending on their country of origin, research topic, etc.). However, they all attended the same event for 4 consecutive days and most of them stayed in the same hotel and attended the same sections (note, though, that Infocom is a multi-track conference).
- 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 University of Cambridge Computer Laboratory, specifically undergraduate year1 and year2 students, and also some PhD and Masters students. This dataset 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 9 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.

3 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 statistical properties for the two conference scenarios are quite similar. To prevent redundancies, the later sections we only selectively show one example, in most cases *Infocom06*, since it contains more participants.

Figure 1 and Figure 2 show the contact duration and number of contacts distribution for each pair in four experiments. For the *HongKong* experiment we include the external devices, but for the other three experiments we use only the internal devices. We show later that for the *HongKong* experiment we need to use the external devices to help to forward the data because of network sparseness.

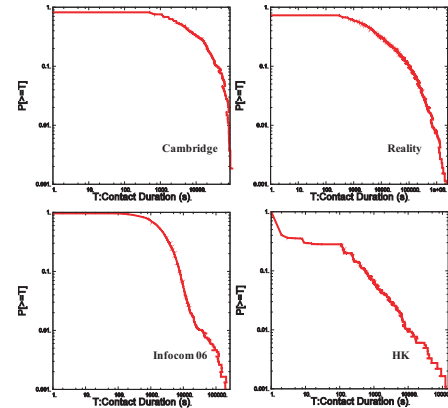


Fig. 1. The contact duration distribution for each pair in four experiments.

3.1 Correlation between Regularity and Familiarity

We assume that contact duration implies *familiarity*. Two people sharing the same office might hate each other, and not talk, but we will ignore this kind of extreme situation here. The number of times two people meet each other implicitly reveals the pattern with which they meet. In this work, we infer *regularity* of meetings from the number of contacts. Two people might meet a lot of times in a short period (e.g. a day), and then not at all. However, short periods with many contacts are less likely to contribute to the upper quarters of the distribution, and here we will ignore these too as outliers.

Figure 3 shows the correlation between Number Contact (Regularity) and Contact Duration (Familiarity) in four the

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

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

TABLE 1
Characteristics of the five experimental data sets




Fig. 2. The number of contacts distribution for each pair in four experiments.

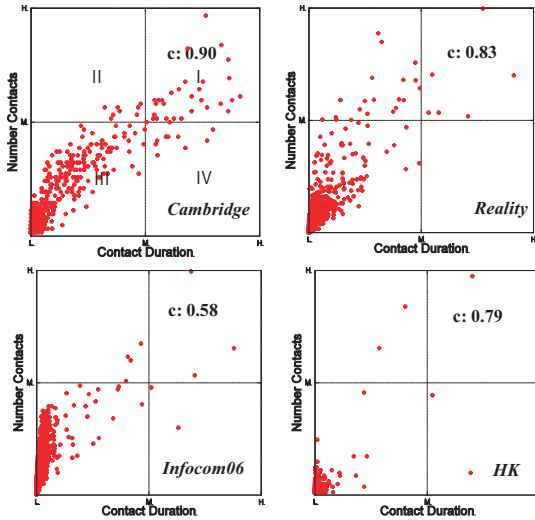


Fig. 3. Number of contacts versus the contact durations for pairs of *Cambridge* Students.

four datasets. In *Cambridge*, the regularity is positively correlated to familiarity with a correlation coefficient of 0.9026. We can define four kinds of relationships between a pair of nodes: Community (I), Familiar Strangers (II), Strangers (III), and Friends (IV). A pair of nodes which has long contact duration (high familiarity) and large number of contacts (high regularity) is likely to belong to the same community. A pair of nodes which meet regularly but do not spend time with each other, could be familiar strangers [13] meeting everyday. People who do not meet regularly and do not spend time with each other would be in the category of strangers. Finally, for node pairs which do not meet very frequently but spend quite

a lot of time together for each meeting, we count as friends. It is not necessary that the division of the four quarters are exactly at the middle. It is just as a reference or example. A clear-cut division may need more empirical experimental results. But here we provide the methodology to classify these four kinds of relationship based on pure contact duration and frequency. Additional difficulties faced by empirical social network research are well described in work by Watts [14]. For the conference scenario, we can see a narrow stripe in the left bottom quarter. This stripe shows that people in the conference tend to meet each other more often than spend long time together. That is a typical conference scenario, since people may meet each other many times in coffee breaks, corridors, the registration desk etc. They may stand together and chat for a while, and then shift to chat with others instead of spending all the time together. The *Reality* set is similar to the *Cambridge* one, with most of the points lying on or above the diagonal line. However, it seems that people also have more contacts instead of spending time together. In the *HongKong* figure, we can find two pairs of friends, two pairs of close community members, and two pairs of familiar strangers. All the other pairs lie in the strangers quarter. This is in line with our expectations for an experiment designed to contain little social correlation.

These correlation graphs give us an overview idea about the social characteristics of the people in these environments. Later when we study community detection, we show how to detect different types of relationships by choosing suitable thresholds for contact duration and number of contacts based on the division into four quadrants in Figure 3.

4 INFERRING HUMAN COMMUNITIES

A social network consists of a set of people forming meaningful relationships, where prominent patterns of meeting or information flow are observed. In PSN, the social network could map to the computer network since people carry the computing devices. To answer the second question we need community detection algorithms. In this section, we introduce and evaluate two centralised community detection algorithms: *K-CLIQUE* by Palla *et al.* and weighted network analysis (WNA) by Newman [15] [16]. We use these two centralised algorithm to uncover the community structures in our mobile traces. We believe our evaluation of these algorithms can be useful for future traces gathered by the research community.

Many centralised 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 centralised detection method are the ability to uncover overlapping communities, and a high degree of automation (low manual involvement). In real 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 doesn't 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 compliment one another.

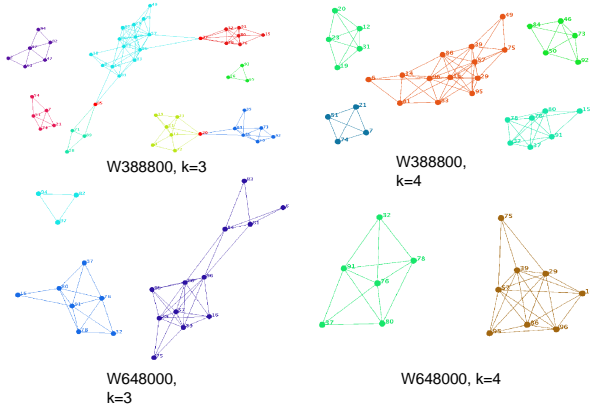


Fig. 4. Communities based on contact durations with weight threshold = 388800s (4.5days), 648000s (7.5days) and $k=3,4$ (*Reality*)

4.1 K-CLIQUE Community Detection

Palla *et al.* 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 datasets above. Here we take *Reality* as an example, since it contains a reasonably large number of nodes, and lasts for a long period of time. Out of 100 experimental participants, 75 are either students or faculty in the MIT Media Laboratory, while the remaining 25 are incoming students at the adjacent MIT Sloan business school. Of the 75 users at the Media lab, 20 are incoming masters students and 5 are incoming MIT freshmen.

First we look at communities detected by using a contact threshold of 388,800 seconds or 4.5 days on the 9 months *Reality* dataset. This threshold was obtained by assuming 3 lectures per week; with 4 weeks per month and a total trace duration of 9 months (2% of the total links are taken into consideration). Research students in the same office may stay together all day, so their contact duration threshold could be very large. For students attending lectures, this estimation can be reasonable. Using a looser threshold still detects the links with much stronger fit. We observe 8 communities of size

(16,7,7,7,6,5,4,3) when $k=3$. When $k=4$, the 3-clique community is eliminated and other communities shrink or are eliminated, and only 5 communities of size (13,7,5,5,4) left. All of these 5 communities are disjoint. When $k=5$, 3 communities of size (9,6,5) remains, the size-9 one and the size-5 one are split from the 13-sized one in the 4-clique case. Moving to $k=6$ and $k=7$, there are 2 communities and 1 community respectively. We are also interested in knowing about small groups which are tightly knit. We set a strict threshold of 648,000 seconds, that is on average 1 hour per weekday, 4 weeks per month, and for a total of 9 months. Around 1% of the links are taken into account for the community detection. When $k=3$, there are three disjoint communities of size (12,7,3). When $k=4$, there are only two communities left of size (8,6). Figure 4 shows the 3-clique and 4-clique communities of 648,000 seconds threshold with its counter parts of 388,800 seconds. A single 7-clique community remains in $k=5$ and $k=6$ cases, this 7-clique community is the same as in the 388,800 second case. These 7 people could be people from a same research group, they know each other and have long contact with each other.

4.2 Weighted Network Analysis

In this section, we implement and apply Newman's weighted network analysis (WNA) for our data analysis [16]. Our contribution is also the extension of the unweighted *modularity* proposed in [17] 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):

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

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 of which 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 summarises the communities detected by applying WNA on the four datasets. 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.

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

TABLE 2
Communities detected from the four datasets

These centralised 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.

5 HETEROGENEITY IN CENTRALITY

In mobility models such as random way point, nodes are often assumed, explicitly or implicitly, to have homogeneous distributions for speed distribution importance and popularity. Our intuition is that the last two assumptions, at least, are not true. 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. Homogeneity might favour different forwarding strategies for PSNs. In contrast, 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. We are not aware of prior work or empirical evidence on using human popularity or node centrality for information dissemination in mobile networks.

A temporal 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 temporal-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 6).
- 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 normalised 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 10).

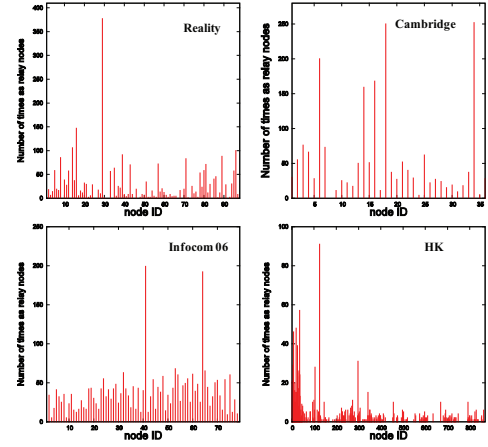


Fig. 5. Number of times a node as relays for others on four datasets.

Figure 5 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 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 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 normalised individual node centrality are shown in Table 3.

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

TABLE 3
Statistics about normalised node centrality in 4 experiments

6 EVALUATION METHODOLOGY

6.1 HaggleSim Emulator

In order to evaluate different forwarding algorithms, we developed an emulator called *HaggleSim*, 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. The event granularity depends on our choice of balance between replay speed, and the degree of accuracy desired. In

all the simulations in this work, we divided the traces into discrete contact events with granularity of 100 seconds. We analyse the successful delivery rate, the delivery cost, the delay distribution, the hop count distribution for all successful deliveries, and the popularity of a node as a relay based on the log files produced by the emulator.

Table 4 shows a snapshot of our emulation source file. 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. As we can see from the source file, some events happen at the same time stamp, and should be treated as simultaneous. Just reading the file line by line artificially imposes an order on events, so instead we keep contacts in a buffer while reading multiple lines ahead: if the next event happens at the same time, and there would be an exchange of messages between the nodes referenced in the line due to the forwarding strategy, we will re-read the contacts in the buffer, and apply the same forwarding strategy to the newly exchanged messages. This removes the artificial ordering.

Node	Node	Time Stamp
6	24	14991
25	11	14991
25	240	14991
24	240	14991
6	240	14991
35	511	14991
24	11	14991
6	24	15001
25	11	15001
25	240	15001

TABLE 4
Snapshot of Simulation Source File

6.2 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.

7 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.

Categories of human contact patterns: Human relationships can be modelled by using correlation of contact duration and number of contacts. We defined four types of human

relationship based on the correlation of contact duration and number of contacts.

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 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 8 to Section 11, we look at how can we use this information to make smart forwarding decisions. The following three pre-existing 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 and cost.

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 datasets. Hence for fair comparison, we would like to evaluate our algorithms against the 4-copy-4-hop MCP scheme in most of the cases.

The Mobile network has a dual nature: it is both a physical network and at the same time it is also a social network. A node in the network is a mobile device, and also associated with a mobile human.

Figure 6 shows the design space for the forwarding algorithms. The vertical axis represents the explicit social structure, that is facets of nodes that can specifically identified such as affiliation, organisation or other social context. 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 localised cohesive structure, and the Structure-in-Degree axis indicates the use of hub structure. These are observable physical characteristics. In our design framework, is not necessary that physical dimensions are orthogonal to the social dimension, but since they are 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 organisation. Optimisations are examined by comparing label of the potential relay nodes

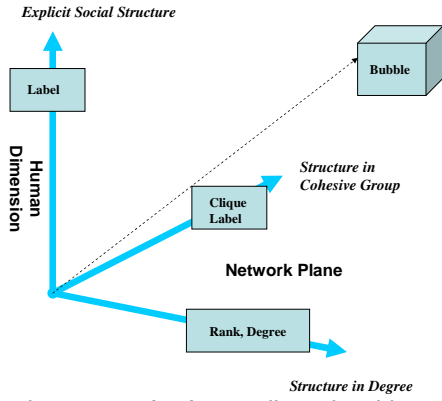


Fig. 6. Design space for forwarding algorithms.

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: This is analogous to the degree of a node in a fixed network; we use a modified ranking scheme, namely the node centrality in a temporal network. 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: A heuristic based on the observed average of the degree of a node over some longer interval. Either the last interval window (S-Window), or a long-term cumulative estimate, (C-Window) is used to provide a fully decentralised 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 with explicit labels, to decide on the best forwarding nodes. BUBBLE is an example algorithm which 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.

In the following sections, we will show how can we make use of these different metrics to improve forwarding performance in a heterogeneous system and also when they will fail.

8 GREEDY RANKING ALGORITHM

The second contribution of this paper is to modify the greedy ranking search scheme over power-law networks to apply to temporal graphs, and evaluate the resulting algorithm.

Here we use a similar greedy strategy to the one Adamic *et al.* introduced in [18]. 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 propose. Here 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. Our strategy, which we call 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. 7. Comparison of delivery ratio (left) and cost (right) of MCP and RANK on 4 copies and 4 hops case (*Reality*).

Figure 7(a) shows that the simple greedy ranking performs almost as well as MCP for delivery. Figure 7(b) also shows that the cost is only around 40% that of MCP, which represents a marked improvement.

Hierarchical organisation is a common feature of many complex systems. The defining feature of a hierarchical organisation is the existence of a hierarchical path connecting any two of its nodes. Trusina *et al.* [19] 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 analysed the percentage of hierarchical paths inside all the shortest paths. Table 5 summarises the results.

The percentage of hierarchical paths is calculated as the number of hierarchical paths divided by the number of non-direct 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

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)

TABLE 5
Hierarchical paths analysis of all shortest paths

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.

RANK appears to work in small and homogeneous systems, but when we look at a more diversified system, for example the *Hong Kong* dataset, 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% 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% 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.

9 DIRECT LABELING 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* dataset [20]. This is a beginning of social based forwarding in PSN, but lack of mechanisms to move messages away from the source when the destinations are socially far away (such as *Reality*). Our third contribution 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* dataset. We can see from Figure 8 that LABEL only achieves around




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

55% of the delivery ratio of the MCP strategy and only 45% of the flooding delivery although the cost is also much lower. 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.

Figure 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% 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% of the first hop nodes (from the D-Group plot) are in the same group as the destination. This explains why direct labeling is not effective, since it is far from discovering the shortest path. We can also see that on going to the 2nd hop, S-Group correlation drops to slightly less than 30%, and when going to 4th-hops, almost all (90%) 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.




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

10 CENTRALITY MEETS COMMUNITY

The fourth contribution in this paper is to combine the knowledge of both centrality 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.

10.1 Two-community Case

In order to make the study more systematic, we start with the two-community case. We use the *Cambridge* dataset for 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 myself to just the two 10-clique groups found with a number-of-contact threshold of 29; that is where each node at least meet another 9 nodes frequently. Some students may skip lectures and cause variations in the results, so this limitation makes our analysis yet more plausible.




Fig. 10. Node centrality in 2 groups in *Cambridge* data

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 Figures 10(a) and 10(b) 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. Figure 11(a) shows the individual node centrality when traffic is created from one group to another. Figure 11(b) shows the correlation of node centrality within an individual group and inter-group centrality. We can see that points lie more or less around the diagonal line. This means that the inter- and intra- group centralities are quite well

correlated. Active nodes in a group are also active nodes for inter-group communication. There are some points on the left hand side of the graph which have low intra-group centrality but moderate inter-group centrality. These are nodes which move across groups. They are not important for intra-group communication but can perform certainly well when we need to move traffic from one group to another.

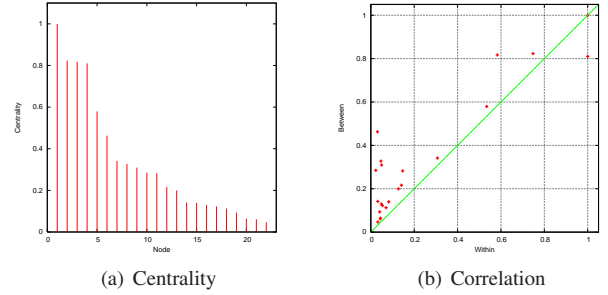


Fig. 11. Inter-group centrality and correlation between intra- and inter-group centrality (*Cambridge*)

We can show now why homogeneous global ranking in Section 8 does not work perfectly. Figure 12 shows the




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

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 push 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-A, since each world/community is like a bubble. Figure 13 illustrates the BUBBLE algorithm.




Fig. 13. Illustration of the BUBBLE forwarding algorithm.

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 reach 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.

We can see from Figure 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% the cost of MCP and BUBBLE-B is even better, with only 45% the cost of MCP. Both have almost the same delivery success as MCP.

10.2 Multiple-community Cases

To study the multiple-community cases, we use the *Reality* dataset. To evaluate the forwarding algorithm, we extract a 3-week session during term time from the whole 9-month dataset. Emulations are run over this dataset with uniformly generated traffic.




Fig. 14. Comparisons of several algorithms on *Cambridge* dataset, delivery and cost.

There is a total of 8 groups within the whole dataset. Figure 15 shows the node centrality in 4 groups, from small-size to medium-size and large-size group. We can see that within each group, almost every node has different centrality.




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




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

In order to make our study easier, we first isolate the largest group in Figure 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 Figure ?? 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 1 week. BUBBLE-A only has 70% and BUBBLE-B only 55% 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; 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 Figure 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 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 2 weeks.² However, the cost is only 50% that of MCP.




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

In order to further justify the significance of social based forwarding, we also compare BUBBLE with a benchmark ‘non-oblivious’ forwarding algorithm, PROPHET[5]. PROPHET uses the history of encounters and transitivity to calculate the probability that a node can deliver a message to a particular destination. Since it has been evaluated against other algorithms before and has the same contact-based nature as BUBBLE (i.e. do not need location information), it is a good target to compare with BUBBLE.

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.




Fig. 18. Comparisons of BUBBLE and PROPHET on *Reality* dataset

Figures 18 shows the comparison of the delivery ratio and delivery cost of BUBBLE and PROPHET. 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 can see that most of the time, BUBBLE achieves a similar delivery ratio to PROPHET, but with only half of the cost. Considering that BUBBLE does not need to keep and update an routing table for each node pairs, the improvement is significant.

Similar significant improvements with BUBBLE are also observed in other datasets. This demonstrates the generality of the BUBBLE algorithm, but for space reasons, we do not include the results here. There are several other ‘non-oblivious’ forwarding algorithms, for example the pattern based Mobyspace Routing [21] by Leguay *et al.*, the location-based routing [22] by Lebrun *et al.*, and adaptive routing by Musolesi *et al.* [23]. Mobyspace and location-based routing both required location related information, but BUBBLE only use contact-based information. They are in different application categories and hence are not directly comparable. Adaptive routing provides a nice framework for choosing relays based on utility calculation, but strictly speaking it is not a forwarding algorithm. There is a social based forwarding algorithm called SimBet routing [24] proposed independently at the same time as BUBBLE. SimBet routing uses ego-centric centrality but without community information to forward data. The RANK algorithm we introduced in this thesis would provide a delivery ratio upper bound for this algorithm. Overall, we evaluated BUBBLE against WAIT, FLOOD, the optimised MCP, LABEL, RANK, and the benchmark PROPHET, it seems sufficient to illustrate the performance of BUBBLE.

11 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

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 dataset.

and calculate its centrality values. In [25], 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% of the centralised K -CLIQUE algorithm. The next step is to ask how can each node know its own centrality in a decentralised way, and how well past centrality can predict the future.

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

11.1 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 Figure 19 that some nodes with a very




Fig. 19. Correlation of rank with total degree and rank with unit time degree (*Reality*).

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 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 sets of emulations using greedy average unit-time degree (or we simply call it DEGREE) instead of the pre-calculated 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 [26] and we would like to do further theoretical investigation.

We can see from Figure 20 that the S-Window approach reflects more recent context, and achieves a maximum of 4% improvement in delivery ratio over DEGREE, 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 DEGREE, which averages throughout the whole experimental period. It does not achieve as good delivery as DEGREE (not more than 10% less in term of delivery), but it also has lower cost.




Fig. 20. Comparisons of delivery (left) and cost (right) of RANK, S-Window and C-Window (*Reality*)

11.2 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 3-week sessions from the *Reality* dataset and then run a set of greedy RANK emulations on the last two data sessions, but using the centrality values from first session.




Fig. 21. Delivery ratio (left) and cost (right) of RANK algorithm on 2nd data session, all groups (*Reality*)

Figure 21 shows the delivery ratio and cost of RANK on the 2nd data session using the centrality values from the 1st 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 dataset which the centrality values derived from. Similar performance is also observed in the 3rd data session. These results imply some 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 decentralised 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 city wide environment, but leave the evaluation details as future work when we can get a larger scale of dataset.

Suppose there is a network of mobile users, perhaps spanning an entire city, each device can detect its own local community using one of the three distributed detection algorithms (e.g. K -CLIQUE). 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.

By combining the methods of estimating centrality from past history (e.g. C-Window), and the distributed community detection algorithms we proposed before [25], we can make BUBBLE work in a distributed way. Or we simply call it DiBuBB algorithm, which uses the distributed K -CLIQUE algorithm to detect local community and C-Window to approximate its own global and local centrality values. Besides that, it operate exactly like BUBBLE.

12 RELATED WORK

For distributed search for nodes and content in power-law networks, Sarshar *et al.* [27] proposed using a probabilistic broadcast approach: sending out a query message to an edge with probability just above the bond³ 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 [28]. Spray and Wait is another “oblivious” flooding scheme but with a self-limited number of copies [29]. Grossglauser *et al.* proposed the two-hop relay schemes to improve the capacity of dense ad hoc networks [30]. 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 Mobyspace Routing by Leguay *et al.* [21], location-based routing by Lebrun *et al.* [22], context-based forwarding by Musolesi *et al.* [23] 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.* [31] 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 [24] uses ego-centric centrality and its social similarity. Messages are forwarded towards the node with higher centrality to increase the possibility of finding the potential carrier to the final destination. In [20]), we use

small labels 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. The evaluation demonstrates that even such a basic approach results in a significant reduction in routing overheads. RANK algorithm introduced in this paper uses betweenness centrality in a similar manner to SimBet routing. On the other hand, BUBBLE exploits further community structures and combines it with RANK for further improvement of forwarding algorithms. The mobility-assisted Island Hopping forwarding [32] 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 emphasise that we take an experimental rather than theoretical approach, which contrasts with other work described above.

13 CONCLUSION

We have shown that it is possible to detect characteristic properties of social grouping in a decentralised fashion from a diverse set of real world human contact traces. We have demonstrated that such characteristics can be effectively used in forwarding decisions. Our algorithms are designed for a delay-tolerant network environment, built out of human-carried devices, and we have shown that they have similar delivery ratio to, but much lower resource utilisation than flooding, control flooding, and PROPHET.

Our decentralised approximation for centrality relates to the predictability of human mobility. We have made an additional contribution in this area by using similarity measures such as the Jaccard index [33]. Further improvement would entail more work on theoretical aspects of graph similarity [34]. We chose 6 hours as basic unit of centrality approximation. This appears to work well on the datasets we used; however, in future work we will examine the sensitivity of the system to this choice of period.

On forwarding, we have not directly emulated the distributed BUBBLE in this paper. Instead, we chose a divide-and-conquer method, showing separately the feasibility of decentralised approximation of centrality due to its inherent predictability. In principle, BUBBLE is supposed to work with a hierarchical community structure, but because of the limited size of data (each experiment is not large enough for us to extract hierarchical structure), the current algorithm and evaluation focus on a flat community structure. This can later be extended to a hierarchical structure. We will further verify our results when more mobility traces are available.

We believe that this approach represents a first step in combining rich multi-level 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 is funded in part by the ITA project, the EU grants for the Huggle project, IST-4-027918, and the SOCIAL-NETS project, 217141. We would like also to acknowledge

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

comments from Steven Hand, Brad Karp, Frank Kelly, Richard Mortier, Pietro Lio, Andrew Moore, Nishanth Sastry, Derek Murray, Sid Chau, Andrea Passarella, and Hamed Haddadi.

REFERENCES

- [1] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proc. WDTN*, 2005.
- [2] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proc. SIGCOMM*, 2003.
- [3] D. Kempe, J. Kleinberg, and A. Kumar, "Connectivity and inference problems for temporal networks," *J. Comput. Syst. Sci.*, 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," in *Proc. WDTN*, 2005.
- [5] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic routing in intermittently connected networks," in *Proc. SAPIR*, 2004.
- [6] S. Okasha, "Altruism, group selection and correlated interaction," *British Journal for the Philosophy of Science*, vol. 56, no. 4, pp. 703–725, December 2005.
- [7] M. E. J. Newman, "Detecting community structure in networks," *Eur. Phys. J. B*, vol. 38, pp. 321–330, 2004.
- [8] L. Danon, J. Duch, A. Diaz-Guilera, and A. Arenas, "Comparing community structure identification," *J. Stat. Mech.*, p. P09008, Oct 2005.
- [9] L. C. Freeman, "A set of measuring centrality based on betweenness," *Sociometry*, vol. 40, pp. 35–41, 1977.
- [10] N. Eagle and A. Pentland, "Reality mining: sensing complex social systems," *Personal and Ubiquitous Computing*, vol. V10, no. 4, pp. 255–268, May 2006.
- [11] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *Proc. INFOCOM*, April 2006.
- [12] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft, "Opportunistic content distribution in an urban setting," in *ACM CHANTS*, 2006, pp. 205–212.
- [13] E. Paulos and E. Goodman, "The familiar stranger: Anxiety, comfort, and play in public places," in *Proc. ACM SIGCHI*, 2004. [Online]. Available: <http://berkeley.intel-research.net/paulos/research/familiarstranger/>
- [14] D. J. Watts, *Small Worlds – The Dynamics of Networks between Order and Randomness*. Princeton, New Jersey: Princeton University Press, 1999.
- [15] G. Palla, I. Derényi, I. Farkas, and T. Vicsek, "Uncovering the overlapping community structure of complex networks in nature and society," *Nature*, vol. 435, no. 7043, pp. 814–818, 2005. [Online]. Available: <http://dx.doi.org/10.1038/nature03607>
- [16] M. E. J. Newman, "Analysis of weighted networks," *Physical Review E*, vol. 70, p. 056131, 2004.
- [17] M. E. J. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Physical Review E*, vol. 69, February 2004. [Online]. Available: <http://arxiv.org/abs/cond-mat/0308217>
- [18] L. A. Adamic, B. A. Huberman, R. M. Lukose, and A. R. Puniyani, "Search in power law networks," *Physical Review E*, vol. 64, pp. 46 135–46 143, October 2001.
- [19] A. Trusina, S. Maslov, P. Minnhagen, and K. Sneppen, "Hierarchy measures in complex networks," *Physical Review Letters*, vol. 92, p. 178702, 2004. [Online]. Available: [doi:10.1103/PhysRevLett.92.178702](https://doi.org/10.1103/PhysRevLett.92.178702)
- [20] P. Hui and J. Crowcroft, "How small labels create big improvements," in *Proc. IEEE ICMAN*, March 2007.
- [21] J. Leguay, T. Friedman, and V. Conan, "Evaluating mobility pattern space routing for DTNs," in *Proc. INFOCOM*, 2006.
- [22] J. Lebrun, C.-N. Chuah, D. Ghosal, and M. Zhang, "Knowledge-based opportunistic forwarding in vehicular wireless ad hoc networks," *IEEE VTC*, vol. 4, pp. 2289–2293, 2005.
- [23] M. Musolesi, S. Hailes, et al., "Adaptive routing for intermittently connected mobile ad hoc networks," in *Proc. WOWMOM*, 2005.
- [24] E. Daly and M. Haahr, "Social network analysis for routing in disconnected delay-tolerant manets," in *Proceedings of ACM MobiHoc*, 2007.
- [25] P. Hui, E. Yoneki, S.-Y. Chan, and J. Crowcroft, "Distributed community detection in delay tolerant networks," in *Sigcomm Workshop MobiArch '07*, August 2007.
- [26] P. Winters, "Forecasting sales by exponentially weighted moving averages," *Management Science*, vol. 6, pp. 324–342, 1960.
- [27] N. Sarshar et al., "Scalable percolation search in power law networks," June 2004. [Online]. Available: <http://arxiv.org/abs/cond-mat/0406152>
- [28] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Duke University, Tech. Rep. CS-200006, April 2000.
- [29] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Spray and wait: An efficient routing scheme for intermittently connected mobile networks," in *Proc. WDTN*, 2005.
- [30] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad-hoc wireless networks," *IEEE/ACM Trans. on Networking*, vol. 10, pp. 477–486, 2002.
- [31] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *Proceedings of the MobiCom 2004*, 2004.
- [32] M. P. N. Sarafijanovic-Djukic and M. Grossglauser, "Island hopping: Efficient mobility-assisted forwarding in partitioned networks," in *IEEE SECON*, 2006.
- [33] P. Jaccard, *Bulletin de la Societe Vaudoise des Sciences Naturelles*, vol. 37, p. 547, 1901.
- [34] V. D. Blondel and P. V. Dooren, "A measure of similarity between graph vertices: Applications to synonym extraction and web searching," *SIAM Rev.*, vol. 46, no. 4, pp. 647–666, July 2004.



Pan Hui is a senior research scientist in Deutsche Telekom Laboratories based in Berlin. He finished his PhD in Computer Laboratory, University of Cambridge in 2007. During his PhD, he also affiliated with Intel Research Cambridge and Thomson Research Laboratory in Paris. Before that he was with University of Hong Kong for his Mphil and bachelor degree. His research interests include delay tolerant networking, mobile networking and systems, planet-scale mobility measurement, social networks,

and the application of complex network science in communication system design. More information about his profile and his research work can be found at <http://www.cl.cam.ac.uk/ph315/>.



Jon Crowcroft is the Marconi Professor of Communications Systems in the Computer Lab, at the University of Cambridge. Jon graduated in Physics from Trinity College, Cambridge University in 1979, and got an MSc in Computing in 1981, and PhD in 1993 both from UCL. He is a fellow of the ACM, the British Computer Society, the IET, the royal academy of engineering and the IEEE. Jon is a recipient of ACM Sigcomm Award in 2009.



Eiko Yoneki is a Post Doctoral Research Associate in the University of Cambridge Computer Laboratory, Systems Research Group, working on the Haggie Project (EU-FP6) and the SOCIALNETs project (EU-FP7). Eiko has received her Ph.D. degree from the University of Cambridge in December, 2006 and a Postgraduate Diploma in Computer Science from the University of Cambridge in 2002. Previously, she has spent several years with IBM (US, Japan, Italy and UK) working on various networking

products. Her research interests are distributed systems over mobile/wireless networks including delay tolerant networks, bio/socio-inspired networks, wireless sensor networks, and event-driven systems. Further details are available at <http://www.cl.cam.ac.uk/ey204>.