Delegation Forwarding

Vijay Erramilli, and Mark Crovella
Dept. of Computer Science
Boston University
evijay@cs.bu.edu, crovella@cs.bu.edu

Augustin Chaintreau, and Christophe Diot
Thomson Paris
Paris, France
augustin.chaintreau@thomson.net, christophe.diot@thomson.net

ABSTRACT

Mobile opportunistic networks are characterized by unpredictable mobility, heterogeneity of contact rates and lack of global information. Successful delivery of messages at low costs and delays in such networks is thus challenging. Most forwarding algorithms avoid the cost associated with flooding the network by forwarding only to nodes that are likely to be good relays, using a quality metric associated with nodes. However it is non-trivial to decide whether an encountered node is a good relay at the moment of encounter. Thus the problem is in part one of online inference of the quality distribution of nodes from sequential samples, and has connections to optimal stopping theory. Based on these observations we develop a new strategy for forwarding, which we refer to as delegation forwarding.

We analyse two variants of delegation forwarding and show that while naive forwarding to high contact rate nodes has cost linear in the population size, the cost of delegation forwarding is proportional to the square root of population size. We then study delegation forwarding with different metrics using real mobility traces and show that delegation forwarding performs as well as previously proposed algorithms at much lower cost. In particular we show that the delegation scheme based on destination contact rate does particularly well.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Store and forward networks*

General Terms

Algorithms, Measurement, Performance

Keywords

Mobile Opportunistic Networks, Delay-Tolerant Networks, Forwarding Algorithms, Optimal Stopping, Pocket Switched Networks

MobiHoc'08, May 26–30, 2008, Hong Kong SAR, China.

1. INTRODUCTION

Mobile opportunistic networks are comprised of human-carried mobile devices moving in a restricted physical space [6]. Examples include individuals moving in conferences, university campuses and in social settings. They are characterized by nodes with heterogeneous contact rates, unpredictable mobility and limited information. Communication in such settings relies on both transport of messages by mobile nodes as well as multi-hop forwarding. An important performance metric in such networks is cost, which we define as the total number of message replicas created.

Forwarding algorithms can be placed on a spectrum from epidemic forwarding [23] which relies on flooding the network with messages to wait-for-destination scheme in which a source node forwards only if it encounters the destination. While the former scheme guarantees delivery of the message if a path exists, it comes at a high cost; the latter scheme has the least cost but also has a low success rate.

Most forwarding algorithms seek to find a middle ground between these two extremes by relying on information that can be learned during contacts. Algorithms differ in the type of information used as well as how it is used; however, many algorithms make use of some kind of forwarding metric. We refer generically to the value of a node's metric as its *quality*. At any contact, a node with a lower quality metric will forward messages to the node with higher quality. Examples of this include FRESH [9], in which a node will forward if it encounters another node that has seen the destination more recently; greedy-total [11] in which a node will forward if it encounters nodes with a higher contact rate than itself and SimBet routing [8] which relies on a metric calculated using social analysis techniques. In essence, most algorithms try to select good intermediate nodes for forwarding using only local information while endeavoring to minimize costs. A secondary goal is to ensure fairness by balancing costs across nodes.

Our work starts from the observation that to reduce costs, we might seek to forward *only* to the *highest*-quality nodes. This suggests that the problem consists of making timely forwarding decisions by observing a sequence of samples. The main contribution of this paper is a new forwarding strategy based on this observation that is explicitly designed to reduce costs while achieving high performance. We refer to this as *delegation forwarding*.

The main idea of delegation forwarding is as follows. We assume each node has an associated quality metric. A node will forward a message only if it encounters another node whose quality metric is greater than any seen by the message so far. We show that despite the simplicity of this strategy, it works surprisingly well. We first show analytically that in a N-node network, delegation forwarding has expected cost $O(\sqrt{N})$ while the naive scheme of forwarding to any higher quality node has expected cost $O(N)$. We then study delegation forwarding on real mobility traces. We find that delegation forwarding shows performance as good as other schemes at a much lower cost. We compare algorithms in terms of performance (rate of successful delivery and mean delivery delay) as well as cost (number of message replicas created). We also look into cost imbalance on a per-node basis of different algorithms. We find that while most delegation strategies do a good job in balancing cost, the delegation strategy with destination contact rate as the metric does very well.

2. RELATED WORK

The kind of networks we study in this paper are special cases of delay-tolerant networks (DTNs). In contrast to DTNs where mobility can be predicted or future information is known [18, 13], we assume no regularity of movement patterns and so our approach is naturally more probabilistic in nature. Likewise VANETs offers more opportunity for movement prediction than in our setting [4, 20, 2]. Our work focuses on individuals carrying mobile devices, where messages are transferred using all available communication opportunities including mobility of intermediate nodes [7, 11, 16].

Theoretical analysis of these networks assume a homogeneous network where all nodes are equally likely to meet the destination of a message [9, 7, 22]. In that case, the performance cost tradeoff is simply determined by the number of nodes that are used for a single message. There has been recent work which considers heterogeneous conditions [12], where the authors show the maximum flow that can be achieved by static routing if global information about nodes schedule is known. Our scheme is different as we do not assume global information, and forwarding decisions are made in an online manner when nodes are met.

Many forwarding algorithms [17, 11, 14, 22, 8, 9] aim to reduce cost while achieving performance (success rate and mean delay) similar to that of epidemic forwarding [23]. Most rely on comparisons between per-node metrics to make forwarding decisions. FRESH [9] relies on a node's last encounter time with the destination to make a forwarding decision. Greedy [11] relies on contact rate with the destination and greedy-total uses total contact rate of a node to make forwarding decisions. In addition there exist schemes where the number of times a message can be replicated is pre-specified. Examples include binary spray and wait [22] where a number of replicas > 1 can be generated and Simbet [8] where only one replica of the message is allowed. We consider all these schemes for comparison.

3. DELEGATION FORWARDING

We work in the following setting: we assume a set of mobile nodes $\mathcal{N}_i \in \mathbb{M}$ with $|\mathbb{M}| = N$. Nodes generate messages over time; each message has a particular source $\sigma \in \mathbb{M}$ and destination $\delta \in \mathbb{M}$. At random times nodes come into contact, meaning that they are capable of exchanging messages. Messages are transmitted in whole from node to node at time instants during node contact intervals, after which both nodes hold message *replicas*. In our analysis we make no assumptions about the time instants when messages are generated or the time needed for transmission; in our simulations we generate messages according to a Poisson process, and messages are transmitted with no transmission lag. Nodes do not possess any *a priori* knowledge of the number of nodes in the system or knowledge of any properties of the other nodes.

The metrics we are concerned with are: (1) *cost,* which is the number of replicas per generated message in the network; (2) *success rate*, which is the fraction of generated messages for which at least one replica is eventually delivered; and (3) *average delay,* which is the average duration between a message's generation and the first arrival of one of its replicas at the destination. By "high performance" we mean high success rate and low average delay. Furthermore, we distinguish per-node cost variants: (1) *node transmission load,* which is the number of message replicas a node has to forward and (2) *node memory load*, which is the number of message replicas a node has to store in its buffer.

For any given node $\mathcal{N}_i$, the forwarding problem in this setting reduces to the simple question: "upon contact with node $\mathcal{N}_j$, which (if any) of the messages I am holding should I forward to $\mathcal{N}_j$?" As stated in Section 1, we abstract the information available during a contact event to a quality metric associated with each node and message. By the nature of the metric, moving the message to a node with higher quality for this message makes the message more likely to be delivered.

For many algorithms, the answer to the forwarding question is "forward message m if $\mathcal{N}_j$'s quality for message m is higher than mine." This is a tradeoff between the high cost of flooding the network and the low success rate of waiting to encounter the destination. However, the cost of this approach can still be quite high, as we show in Section 4.

Algorithm 1 Delegation Forwarding

```
Let N_1, ..., N_N be nodes
Let M_1, ..., M_M be messages
Node N_i has quality x_im and threshold τ_im for M_m.
INITIALIZE ∀i, m : τ_im ← x_im
On contact between N_i and node N_j:
for m in 1, ..., M do
  if M_m is currently held by N_i then
    if τ_im < x_jm then
      τ_im ← x_jm
      if N_j does not have M_m then
        forward M_m from N_i to N_j
      end if
    end if
  end if
end for
```

To reduce costs even more, we make the requirement for forwarding more stringent. Our approach seeks to forward the message only to the *highest* quality nodes in the system. Conceptually, we would like to create a small number of replica copies, and place them with the nodes which are the very best candidates for eventual delivery to the destination. Thus the forwarding question in our approach becomes "is $\mathcal{N}_j$ among the very highest quality nodes for message m?"

Since there is no *a priori* or global knowledge of node quality, our forwarding question is an instance of an *optimal* *stopping problem* [21]. The problem of optimal stopping is concerned with choosing a time to take an action based on sequentially observed random variables in order to maximize an expected payoff; the classic secretary problem is the best-known example [1]. In our case the expected payoff is the average quality of the nodes that eventually are holding the message.

Optimal stopping theory suggests that a simple strategy appropriate for this problem is to select the *maximum over the observations so far*. In fact, this approach has similarities to the hiring strategies studied in [3]. That paper considered a company interviewing candidates one by one and seeking to maximize average employee quality. A possible strategy would be to only hire candidates better than all current employees, called the **max** strategy. In fact, **max** is not a particular good strategy for hiring, since new hires happen more and more rarely as time goes on, as shown in [3]. However, the forwarding problem is subtly different: when messages are forwarded, they are replicated and so become more numerous. This counteracts the slowdown in forwarding rate of any given message replica, and makes **max** surprisingly effective.

Delegation forwarding then consists of using the **max** strategy over quality to answer the forwarding question. A formal statement of delegation forwarding is given in Algorithm 1.

4. ANALYSIS

In this section we analyze delegation forwarding and show that it reduces expected costs dramatically compared to the naive alternative.

4.1 Assumptions

As described in the previous section, we assume the existence of a quality metric with the property that nodes with higher quality are better candidates as intermediate carriers of a message than are nodes of lower quality. The quality metric can be destination-specific or destination-independent. A destination-specific quality metric is one that varies depending on the destination of a message. For example, FRESH [9] uses the time elapsed since the last contact with the destination as a metric. Other quality metrics, such as the total contact rate of a node (used in [11]) is the same for all destinations and hence destination-independent. The results we present on the cost of delegation forwarding apply equally to both cases; hence for simplicity we drop the message subscript m from x_{im} and refer only to x_i. We study cost imbalance only for the destination-independent case. That case is intuitively the worst case scenario for imbalance, since in that case high quality nodes are the same regardless of destination. In Section 5 we will consider forwarding algorithms in simulation having destination-dependent quality metrics and we will see resulting improvements in cost imbalance.

Depending on the forwarding algorithm used, a node's quality metric and contact rate may be dependent or may be independent. For example, choosing the total contact rate of a node as its quality metric has been recently advocated [11]. We show below that our results apply in the two extreme cases where quality and rate of nodes are independent, and when they are identical.

We assume that the node quality metric x_i follows a uniform distribution on the interval $(0, 1]$. On one hand, we note that if the quality metric is the node's contact rate, it corresponds to the distribution observed empirically in some conference settings (see [11, 6]). Note that a uniform distribution of rates is not inconsistent with a power-law distribution of inter-contact times, which have been observed in [5, 6, 15]. On the other hand, we note that when quality is *independent* of contact rate, this assumption is not a restriction. In that case, the absolute value of the metric can be changed arbitrarily as long as the ordering between nodes is preserved. Any distribution can then be mapped to the uniform case.

We further assume that the node quality metric x_i is time-invariant. However we relax this assumption in our simulations, where some of the node quality metrics we use are continously updated.

4.2 Cost

In the following, we consider a single message and study how many times it is forwarded before reaching a node with a high quality metric. This allows us to prove a bound on the number of copies created for each message. The initial part of our analysis makes use of the framework set out in [3].

4.2.1 Quality Independent of Contact Rate

For any node i maintaining a quality metric x_i(which lies between $(0; 1]$ and a threshold value τ_i, we focus on the gap $g_i = 1 - \tau_i$ between the current threshold and 1. The node that generates the message has threshold initially equal to its quality, *i.e.,* $\tau_i = x_i$. We denote the initial gap $g = 1 - x_i$.

Consider a node that updated its gap value n times. We denote the node's current gap as the random variable G_n. Since nodes meet according to rates that are independent of node quality, the node is equally likely to meet a node with any particular quality value. The next update of the gap of the nodes then occurs as soon as it meets a node with a quality greater than G_n, and all values above this threshold are equally likely.

Hence, we can write

$$G_{n+1} = G_n \cdot U\,, \tag{1}$$

where U is independent of G_n and follows a uniform distribution on $(0, 1]$. By induction we then find:

$$\mathbb{E}\left[G_{n+1} \middle| \, G_n\right] = \frac{G_n}{2}\,, \quad \text{hence} \;\; \mathbb{E}\left[G_n\right] = \frac{g}{2^n}\,.$$

Moreover, from Eq.(1), we see that G_n approximately follows a lognormal distribution (see §2.2 in [3]), with median $\frac{g}{e^n}$. Hence the distribution is highly skewed with most of the probability mass below the mean, and so with large probability we have $G_n \leq \frac{g}{2^n}$.

By setting $\mathbb{E}\left[G_n\right] = \epsilon$, we find that the number of handoff stages it will take to get within ϵ of the highest contact rate node is

$$g/\epsilon = 2^n, \quad \text{so} \quad n = log_2(g/\epsilon). \tag{2}$$

If we want to get to the highest contact rate node, then $\epsilon = 1/N$.

Let us describe the replication process via a dynamic binary tree T, which contains all the nodes that have a copy of the message. Initially T contains a single node with associated gap g. Each time a node with a copy of the message meets another node having higher quality than any node seen so far, we create two children of the node. The children represent each of the two nodes, and both have associated the updated gap value. Note that different branches of this tree may grow more quickly than others. We wish to bound the total size of this tree.

We define the set $B = \{i \,|\, x_i \geq 1 - \frac{g}{\sqrt{N}}\}$, which we call the *target set.* We will also identify a subtree of the tree T in which children are excluded for nodes having a threshold above $1 - \frac{g}{\sqrt{N}}$. We call this subtree the *target-stopped tree.*

The essential observation is the following: if n is close to $\log_2(\sqrt{N})$, then except with a small probability, a node at generation n in the tree has a gap at most $g/2^n \leq g/\sqrt{N}$. This is because of the highly skewed nature of the distribution of G_n, as described above. Hence, we can safely assume that the target-stopped tree has depth at most n. Note that the total number nodes of appearing at generations $0, 1, \ldots, n-1$ is at most $2^n = \sqrt{N}$.

We can now bound the size of the entire tree T, since all nodes at generations $n, n+1, \ldots$ are included in the target set B. Hence, the total size of this tree is at most:

$$C_{\text{delegation}}(g) \lesssim \sqrt{N} + |B| = (1 + \sqrt{g}) \cdot \sqrt{N},$$

hence

$$\mathbb{E}\left[C_{\text{delegation}}\right] \lesssim \frac{5}{3}\sqrt{N}\,. \tag{3}$$

In contrast, the usual style of forwarding algorithm makes no threshold adaptation. A message starting at a node with gap g will eventually reach each of the nodes with higher quality, so that the cost

$$C_{\text{no-delegation}}(g) = gN, \quad \text{hence} \quad \mathbb{E}\left[C_{\text{no-delegation}}\right] = \frac{N}{2}\,. \tag{4}$$

Hence we see that delegation forwarding narrows the set of targeted nodes as additional message copies get created. This saves a significant fraction of the cost, while still causing the message to reach the most important nodes.

4.2.2 Quality Equal to Contact Rate

So far we have analyzed the case where node contact rates are independent of their quality metric. In reality, this may not always be the case since in some cases a good candidate for forwarding a message is a node that is met frequently by other nodes. To address this, we consider here an extreme case, namely, where quality and contact rates are identical. This corresponds to a forwarding strategy "forward to high contact rate nodes." We will show that the delegation scheme can take advantage of this correlation, and that the resulting costs are as good or better than in the previous section.

We assume that the contacts between the nodes follow a *product form:* each node i has a total contact rate λ_i, and the rate of contact between nodes i and j is simply equal to the product $\lambda_i \lambda_j / \sum_j \lambda_j$. In other words, the rate of contact for a given pair of nodes depends on the nodes chosen via their total contact rates.

Since quality and contact rates are identical, $x_i = \lambda_i$ and both are distributed uniformly on the interval $(0, 1]$. As a result, we note that nodes with higher quality are met more often than nodes with lower quality. Hence, the quality of the next node met is not distributed uniformly. Instead we have:

$$\begin{aligned}&\mathbb{P}\left[\text{next node met has quality} \in [x, x+dx]\right] = \\ &\quad \mathbb{P}\left[\text{next node met has rate} \in [\lambda, \lambda + d\lambda]\right] = \\ &\quad \frac{\lambda\, d\lambda}{\int_0^1 \lambda\, d\lambda} = 2\lambda\, d\lambda = 2x\, dx.\end{aligned} \tag{5}$$

As before, instead of considering a node's threshold, we consider the gap $g_i = 1 - \tau_i$. Looking at one node, we denote by G_n its gap after n updates. After conditioning on the current value of the gap G_n and substituting $1-g$ for x, we have in expectation:

$$\begin{aligned}\mathbb{E}\left[G_{n+1} |\, G_n\right] &= \frac{\int_0^{G_n} g \cdot 2(1-g)dg}{\int_0^{G_n} 2(1-g)dg} \\ &= G_n \left(\frac{1 - \frac{2}{3}G_n}{2 - G_n}\right).\end{aligned} \tag{6}$$

Note that the function $h : x \to \frac{1-(2/3)x}{2-x}$ is strictly decreasing, approaching $\frac{1}{2}$ when $x \to 0$ and $\frac{1}{3}$ when $x \to 1$. Hence we have,

$$\frac{G_{n-1}}{3} < \mathbb{E}\left[G_n | G_{n-1}\right] < \frac{G_{n-1}}{2},$$

and so by induction,

$$\frac{g}{3^n} < \mathbb{E}\left[G_n\right] < \frac{g}{2^n}.$$

So we obtain a gap on average within ϵ of the minimum value after n handoffs, with n in the range

$$\log_3(g/\epsilon) < n < \log_2(g/\epsilon).$$

This shows the number of handoff stages needed to get within ϵ of the highest quality node is less than in the independent case. Hence the independent case expression $C_{\text{delegation}}(g) \lesssim (1 + \sqrt{g})\sqrt{N}$ is an upper bound for the delegation cost in the case where quality is equal to contact rate, and so the improvement in cost when using delegation forwarding is even greater in this case.

4.2.3 Numerical Results

To confirm our analytical results, we simulated a collection of nodes interacting randomly with a mixture of contact rates, and generating messages uniformly at random with respect to source and destination, where the quality metric used was total contact rate. In Fig. 1 we compare the measured cost (message copies per message created) in simulation to that predicted by Equations (3) and (4). The plot shows that in practice the costs of delegation forwarding as a function of number of nodes is close to predictions, and it confirms the dramatic improvement in cost when using delegation forwarding.

4.3 Cost Imbalance

While the last section showed that the overall cost in terms of message replicas is dramatically reduced under delegation forwarding, it is also important to ask whether the costs are fairly (i.e., equally) shared among the nodes.

To answer this question, we proceed in stages. We return to making the assumption that a node's quality and contact rate are independent. Then, the first question we ask is as follows. Given a node $\mathcal{N}$ with quality a that is holding a message. What is the probability that this node will at

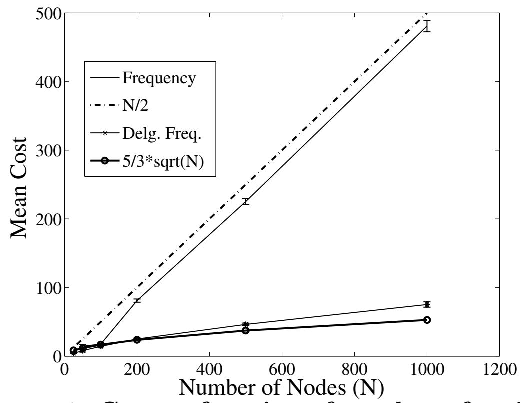

Figure 1: Cost as function of number of nodes.

some point forward this packet to a node $\mathcal{N}'$ having quality x?

Denote this probability density $p(x|a)$. Note that $p(x|a)$ is not a distribution over x because $\mathcal{N}$ may forward its message more than once. However the probability of forwarding to any single node is not greater than 1. That is, for any integer $u < N$, $1/N \cdot p(u/N|0) \leq 1$.

THEOREM 1. *Given a node $\mathcal{N}$ with quality a that is holding a message. The probability that this node will at some point forward this packet to a node $\mathcal{N}'$ having quality $x > a$ is proportional to $\frac{1-a}{1-x}$, for $x \leq 1 - 1/N$.*

PROOF. First consider the case where $a = 0$. Denote the quality of the node forwarded to on forward number i as X_i. The probability distribution of X_i is $p_{X_i}(\cdot)$. Then $p(x|0)$ is equal to:

$$
\begin{aligned}
p(x|0) &= \sum_{i=1}^{\infty} p\begin{pmatrix}\mathcal{N} \text{ forwards to } \mathcal{N}' \\ \text{on forward } i \text{ and not before}\end{pmatrix} \\
&= \sum_{i=1}^{\infty} p_{X_i}(x) \qquad (7)
\end{aligned}
$$

These events are mutually exclusive so the sum is valid. Now the distribution of X_1 is the uniform distribution on (0,1]. We can obtain X_i for successive is by the law of total probability:

$$p_{X_i}(x) = \int_0^1 p_{X_i}(x|X_{i-1} = u) p_{X_{i-1}}(u)\, du$$

$$
\begin{aligned}
p_{X_2}(x) &= \int_0^1 p_{X_2}(x|X_1 = x_1) p_{X_1}(x_1)\, dx_1 \\
&= \int_0^x 1/(1 - x_1)\, dx_1 \\
&= -\ln(1 - x)
\end{aligned}
$$

In the same manner we find that $p_{X_3}(x) = \frac{1}{2}\ln^2(1 - x)$ and $p_{X_4}(x) = -\frac{1}{6}\ln^3(1 - x)$. So we can see that $p_{X_n}(x) = \pm\frac{1}{n-1!}\ln^{n-1}(1 - x)$ with alternating signs. Returning to Equation (7), we can now write:

$$
\begin{aligned}
p(x|0) &= \sum_{i=1}^{\infty} p_{X_i}(x) \\
&= 1 - \ln(1 - x) + \frac{1}{2}\ln^2(1 - x) \\
&\quad - \frac{1}{6}\ln^3(1 - x) + \ldots
\end{aligned}
$$

To evaluate this sum, consider it a function of x. That is, $f(x) = p(x|0)$. Then we note from differentiating the infinite sum above that $f'(x) = \frac{1}{1-x} f(x)$. Thus

$$f(x) = \frac{k}{1 - x} \quad \text{for some } k > 0$$

since then $f'(x) = k/(1-x)^2 = 1/(1-x)\, f(x)$. Since $f(0) = 1$, we find that $k = 1$.

For the case when $a > 0$, we reason that the same relationship should apply, with the range $(a, 1]$ mapped to the range $(0, 1]$. Then,

$$p(x|a) = \frac{1}{1 - \frac{x-a}{1-a}} = \frac{1 - a}{1 - x}$$

Where the expression $\frac{x-a}{1-a}$ maps an x in the range $(a, 1]$ to the range $(0, 1]$. □

This agrees with intuition: if a node with quality zero is holding the message, the probability that a node of quality $1/N$ will receive the message is $1/N$; a node of quality $1/2$ will receive the message with probability $2/N$; and a node of quality $1 - 1/N$ will receive the message with probability 1.

Now we can begin to answer the question of cost imbalance. Given a node $\mathcal{N}$ having quality x, what is the expected number of messages M that will be sent to it?

Denote the probability that a node with quality x generates a message as $p_m(x)$. Assume that messages are generated uniformly, so $p_m(x) = 1$. Then this is

$$
\begin{aligned}
E[M|x] &= \int_0^x p(x|a)\, p_m(a)\, da \\
&= \int_0^x \frac{1-a}{1-x}\, da \\
&= \frac{2x - x^2}{2 - 2x} \qquad (8)
\end{aligned}
$$

For example, assume each node generates one message and there are 100 nodes. Then the graph in Figure 2(a), which is a plot of Equation (8) at intervals of 1/100, shows how many messages each node should be expected to receive. For example, node 99 will receive 50 messages.

Further, we can also address the question of forwarding cost. Given the same assumptions as above, how many messages should a node expect to have to forward, i.e., what is $E[F|x]$?

We reason, as in the analysis of overall forwarding cost in Section 4.2.1, that a node having quality x and holding a message needs to forward it approximately $k(x)$ times, where $k(x)$ is given by:

$$2^{k(x)} = \frac{(1 - x)N}{2^{k(x)}}.$$

As before, the reasoning behind this is that it is approximately the case that at some point after n forwards, 2^n nodes will be holding the message, and there will be $gN/2^n$ nodes of higher quality not holding the packet. Here g is the initial gap, i.e., $1 - x$. Then there will be $n + 1$ forwards altogether.

This means $k(x)$ is $n + 1 = \log_2((1 - x)N)$. So to answer our question,

$$E[F|x] = E[M|x]k(x).$$

Continuing the same example, the expected number of forwarded messages is shown in Figure 2(b). The median value is 4.2 and the maximum is 25.

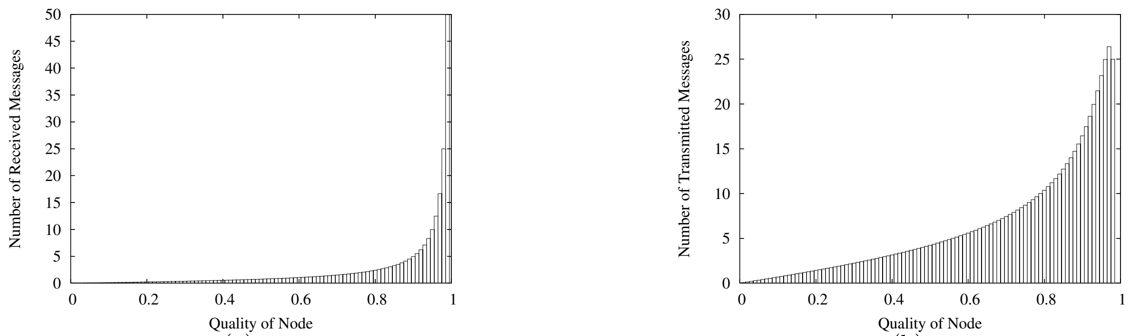

(a) (b)

Figure 2: Cost Imbalance, 100 nodes. (a) Node Memory Load and (b) Node Transmission Load; as a function of quality.

Thus we conclude that both node memory load and node transmission load are unevenly distributed. Node memory load is more highly skewed than node transmission load because as a node's quality increases, the number of nodes it must send to diminishes. In a 100 node system the busiest node has about five times the transmission load of a typical node, and the top 5% have at least four times the transmission load of a typical node. Overall, this level of imbalance is undesirable but not prohibitive.

5. PERFORMANCE ON REAL TRACES

In the previous section we analyzed delegation forwarding and showed that it can dramatically reduce costs. In this section we show that, given realistic contact patterns, delegation forwarding can yield performance comparable to non-delegation approaches, and comparable to the best known forwarding algorithms. Further, we assess the degree of cost imbalance under delegation forwarding and show that in realistic contact patterns, certain delegation forwarding algorithms show cost imbalance no worse than the best alternatives.

5.1 Data Details

In order to investigate delegation forwarding schemes in realistic settings, we use a diverse collection of empirical data sets.

Contact Traces: The first group of data sets consists of contact traces between short-range Bluetooth enabled devices (iMotes [7]) carried by individuals in conference environments, specifically Infocom 2006 and Conext 2006. We isolated two 3-hour periods from both the data sets for our study. We selected the 3-hour periods such that the total contact rate of nodes is fairly stable.

AP-Based Traces: The second group of data sets consist of two groups. First, we use contact traces gleaned from a university campus (UCSD) [19]. The data set consists of of client based logs of WiFi access points. Second, we use traces from the MIT RealityMining [10] project. These datasets include Bluetooth device-device contact data between 100 users as well, as connection logs from GSM towers. Information on how device-device contact traces were obtained as well as more details on the data sets can be found in [7]. For the UCSD dataset, we selected two apparently-stable 6-hour periods and for reality mining we selected two apparently-stable 3 hour periods (one for device-device logs and one for GSM based).

5.2 Experiments

We implemented a variety of forwarding algorithms in a trace-driven simulator. For each trace and forwarding algorithm we study, we generate a set of messages with sources and destinations chosen uniformly at random, and generation times from a Poisson process averaging one message per 4 seconds. We assume nodes have infinite buffers and carry all message replicas they receive until the end of the simulation. All our results are averaged over 10 simulation runs. Our metrics are success rate, average delay, and cost (as defined in Section 3). Each simulation run was therefore 3 or 6 hours depending on the trace; to avoid end-effects, no messages were generated in the last hour of each trace.

We selected forwarding algorithms so as to include both well-known existing algorithms as well as algorithms that span a wide range of design choices. All these algorithms are distributed and operate in an online manner working with local information. For each algorithm we describe the rule used to decide whether to forward a message $\mathcal{M}_m$ held by $\mathcal{N}_i$ when node $\mathcal{N}_i$ meets node $\mathcal{N}_j$.

Epidemic (Flooding) [23]: Node $\mathcal{N}_i$ forwards $\mathcal{M}_m$ to $\mathcal{N}_j$ unless $\mathcal{N}_j$ already has a replica of $\mathcal{M}_m$. Epidemic forwarding achieves the best possible performance, so this algorithm yields upper bounds on success rate and average delay. However it is also the case that epidemic forwarding will have the highest costs.

Frequency [11]: Node $\mathcal{N}_i$ forwards $\mathcal{M}_m$ to node $\mathcal{N}_j$ if $\mathcal{N}_j$ has more total contacts (with all other nodes) than does $\mathcal{N}_i$. This algorithm is destination independent. This is referred to as the greedy-total scheme in [11].

Last Contact: Node $\mathcal{N}_i$ forwards $\mathcal{M}_m$ to $\mathcal{N}_j$ if $\mathcal{N}_j$ has contacted any node more recently than has $\mathcal{N}_i$. This algorithm too is destination independent.

Destination Frequency: Node $\mathcal{N}_i$ forwards $\mathcal{M}_m$ to $\mathcal{N}_j$ if $\mathcal{N}_j$ has contacted $\mathcal{M}_m$'s destination more often than has $\mathcal{N}_i$.

Destination Last Contact [9]: Node $\mathcal{N}_i$ forwards $\mathcal{M}_m$ to $\mathcal{N}_j$ if $\mathcal{N}_j$ has contacted $\mathcal{M}_m$'s destination more recently than has $\mathcal{N}_i$. This algorithm is also known as FRESH [9].

Spray and Wait (SpWt) [22]: $\mathcal{M}_m$'s source initially creates l replicas of $\mathcal{M}_m$. If node $\mathcal{N}_i$ has $k > 1$ replicas of $\mathcal{M}_m$ and $\mathcal{N}_j$ has no replicas, $\mathcal{N}_i$ will forward half its replicas to $\mathcal{N}_j$ and keep the other half. If node $\mathcal{N}_i$ has just one replica of $\mathcal{M}_m$, it uses the destination last contact rule (described above). We use two variants, having $l = 4$ and 8.

SimBet [8]: Node $\mathcal{N}_i$ forwards $\mathcal{M}_m$ to node $\mathcal{N}_j$ if $\mathcal{N}_j$

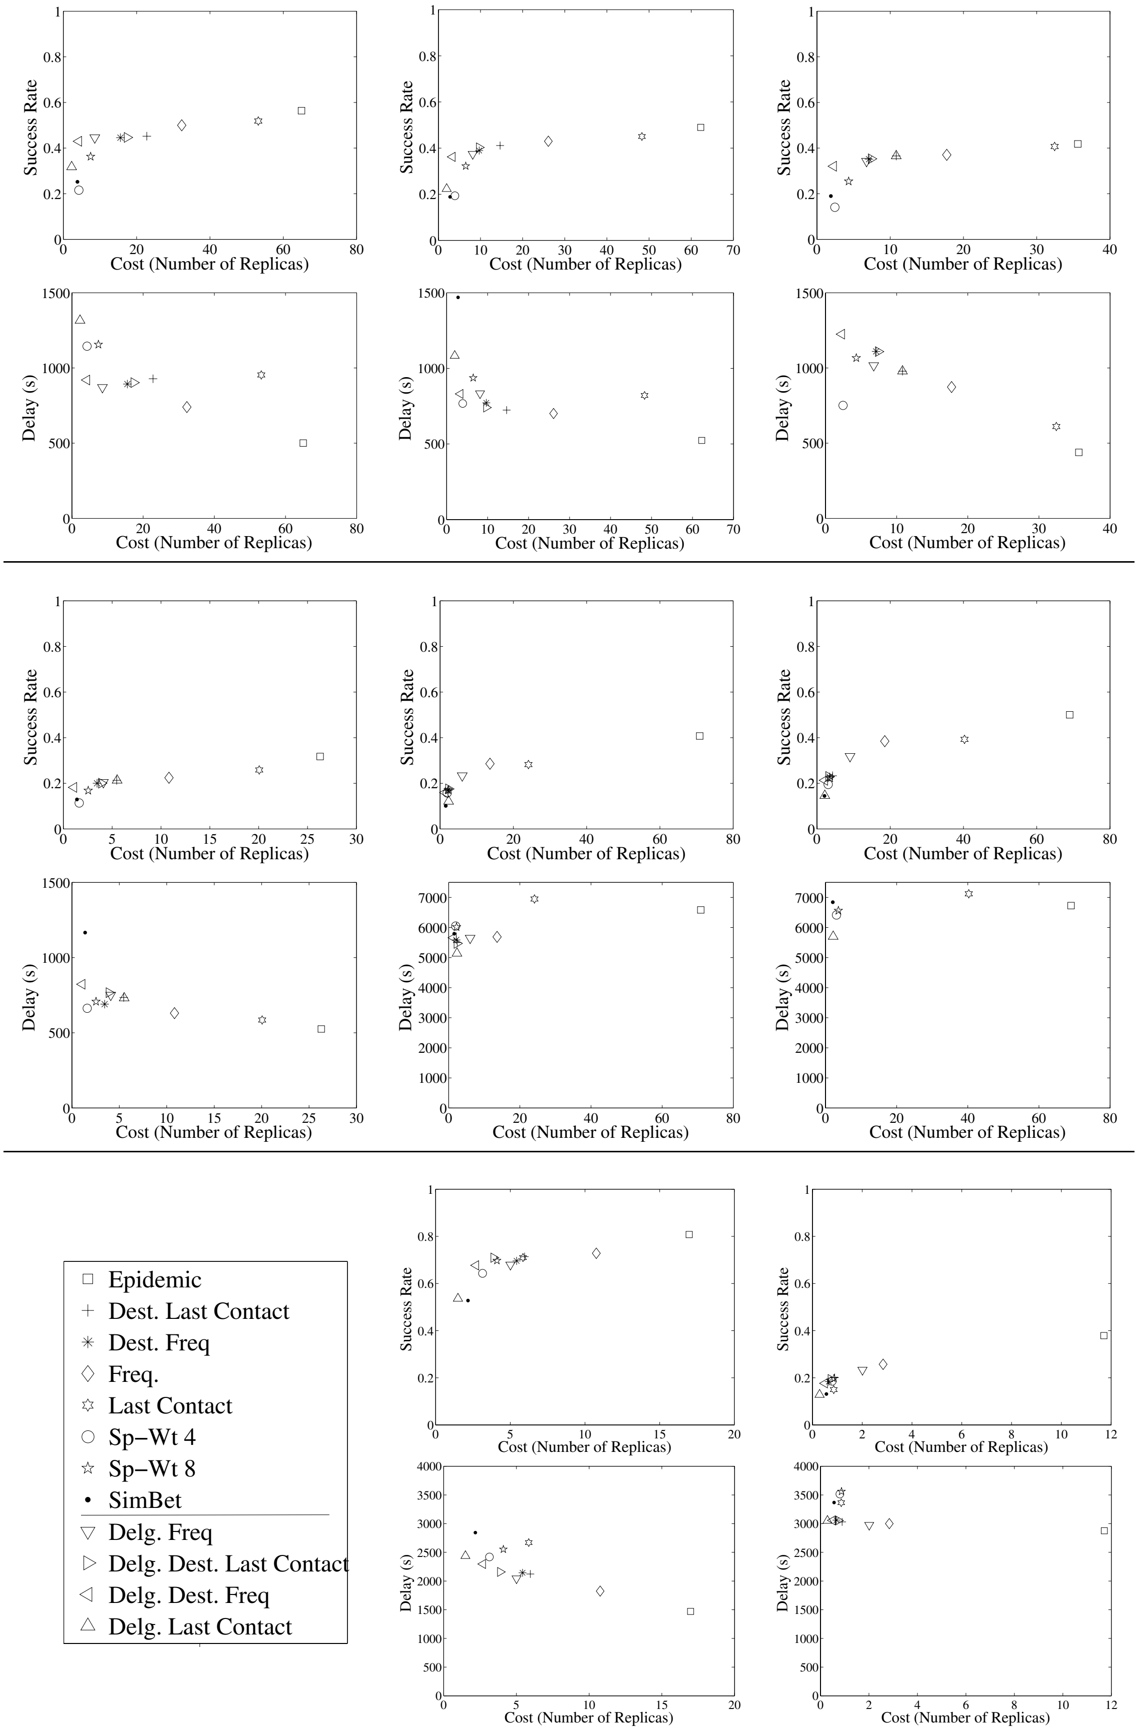

Figure 3: Performance of Forwarding Algorithms. Top Row: Infocom 06 9-12; Infocom 06 3-6; Conext 06 9-12. Middle Row: Conext 06 3-6; UCSD 9-3; UCSD 3-9. Bottom Row: RealityGSM; RealityBT. For each dataset, upper plot is Cost versus Success Rate; lower plot is Cost versus Mean Delay.

scores higher on the simbet metric. To compute the simbet metric, one views the underlying contact graph as a social encounter graph, and incorporates two social measures (similarity and betweenness) of a node. Only one replica of the message exists in the network. We use the same parameters used in [8].

The delegation schemes we consider are: **Delegation Destination Frequency**, **Delegation Destination Last Contact**, **Delegation Frequency**, and **Delegation Last Contact**, each of which is obtained by applying the delegation forwarding strategy to the corresponding algorithm above.

5.3 Results

Cost.

The results of our simulations are shown in Figure 3. Each plot shows either success rate (upper) or mean delay (lower) versus cost. The eight pairs of plots correspond to the eight traces we analyzed, as listed in the caption to the figure. In these plots, the four triangles (at different orientations) correspond to the four delegation forwarding algorithms we studied.

In the success rate versus cost plots, the best algorithms are those closest to the upper left corner; in the delay versus cost plots the best are in the lower left corner. Our first observation is that usually, one of the delegation algorithms occupies the best position in these plots. This means two things: first, as expected, delegation forwarding has very low costs – usually the lowest of any algorithm. Second, and more surprisingly, delegation forwarding usually performs about as well as most other forwarding algorithms. Indeed, delegation schemes reduce cost drastically, by as much as 3/4 of the original cost, while maintaining approximately the same success rate with a modest increase in average delay. Taken as a whole across all these plots, delegation forwarding is clearly the best choice for trading off performance and cost.

We note that in terms of success rate, it is often the case that most forwarding algorithms have performance within a narrow range. This is also true for mean delay, although the range of variation can be larger. These results suggest that forwarding algorithms differ mostly in the costs incurred, rather than performance, which is consistent with previous results [11] and argues for favoring delegation forwarding.

Looking more closely at the variants of delegation forwarding, we see that delegation destination frequency is most often the best of the delegation approaches. This makes sense because, being destination-specific, the algorithm has more information to work with, and hence becomes more selective in forwarding.

Cost Imbalance.

In Section 4 we showed analytically that node memory and node transmission loads can be unevenly distributed across nodes when the quality metric is independent of destination. It is important to ask how serious this problem is in realistic settings, and also whether destination-dependent quality metrics show different behavior.

To answer these questions we plot the node transmission load and node memory load (normalized) for three non-delegation approaches and two variants of delegation forwarding. The results are shown in Fig. 4. Fig. 4(a) shows that all algorithms suffer approximately equally from imbalance in node transmission load, except delegation frequency which has noticeably higher imbalance. The imbalance seen for delegation frequency (where quality of node is equal to its contact rate) is in line with analytical results (Section 4), e.g., the top 5% of the nodes transmit almost 5 times more than a typical node. Delegation destination frequency however shows no more imbalance than any of the non-delegation schemes. This occurs because under the destination frequency approach, the highest quality nodes for different destinations tend to be different. (We have verified this via additional analysis, but space does not permit their inclusion.)

Fig. 4(b) shows imbalance in node memory load. As expected from analysis, node memory load is more skewed for delegation frequency than was node transmission load. For example, the top 10% of the nodes have around 8 times the load of a typical node, consistent with analysis. Once again however, delegation destination forwarding shows no greater imbalance than the non-destination approaches.

6. CONCLUSIONS

Forwarding in mobile opportunistic networks is a challenging problem. An objective of many forwarding algorithms is to reduce cost while trying to keep success rate high and delays low. These algorithms often rely on a quality metric associated with nodes in the network to make forwarding decisions. The main contribution of this paper is to propose a new forwarding strategy explicitly designed to reduce costs while maintaining the high performance of such algorithms.

We start with the observation that in order to reduce costs, a good strategy could be to forward *only* to the *highest*-quality nodes. However the distribution of node quality is not known, so every node has to decide if a current encounter is with a high-quality node or not by relying on past contacts. This is an instance of an optimal stopping problem, and seen in that light, a good strategy is for a node to only forward a message to nodes with quality greater than any seen so far for this message. This is delegation forwarding.

We analyze two variants of delegation forwarding, considering both the cases where quality is independent of the underlying contact rate and also where quality is identical to a node's contact rate. In both these cases we show that delegation forwarding reduces expected costs dramatically (from $O(N)$ to $O(\sqrt{N})$) while still ensuring that messages reach the highest quality nodes. We also study the fairness of delegation forwarding in terms of the distribution of per-node cost. We show that when node quality is independent of message destination, delegation forwarding can induce a moderate level of imbalance in per-node cost.

We then turn to studying the performance of delegation forwarding on real mobility traces and observe that overall, delegation forwarding approaches are preferable to any of a set of commonly-studied forwarding algorithms. Delegation approaches generally achieve comparable performance to well-known alternatives while also generally achieving remarkably low costs. Furthermore, we show that when evaluated in realistic settings, delegation forwarding using destination contact frequency as its metric shows no greater per-node imbalance than non-delegation alternatives, suggesting that the worst-case results from analysis need not be experienced in practice.

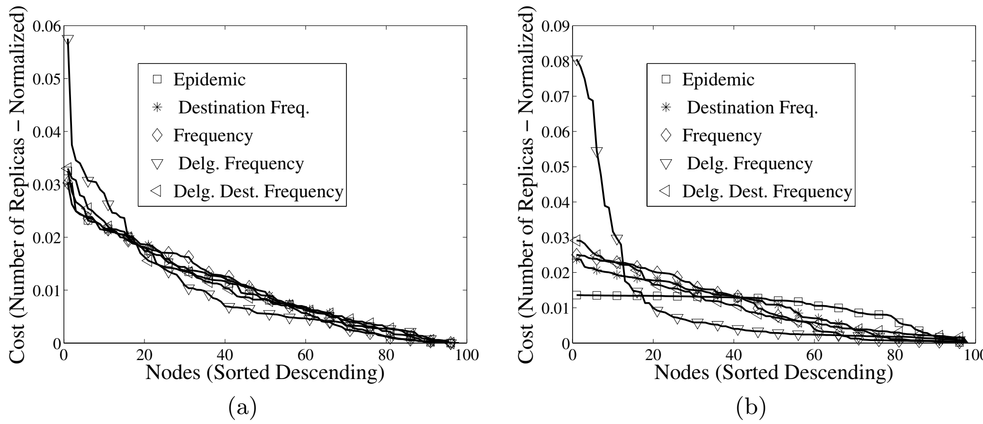

Figure 4: Plots (Infocom 9-12) (a) Number of Transmissions (b) Queue Lengths

Acknowledgments

The authors thank John Byers, Ibrahim Matta and Fernando Silveira for helpful discussions. This work was supported by NSF grants CCR-0325701 and ANI-0322990, and by funding from the HAGGLE FET project.

7. REFERENCES

[1] Ajtai, M., Megiddo, N., and Waarts, O. Improved Algorithms and Analysis for Secretary Problems and Generalizations. In *IEEE FOCS* (1995), pp. 473–482.

[2] Balasubramanian, A., Levine, B., and Venkataramani, A. DTN Routing as a Resource Allocation Problem. In *ACM SIGCOMM '07*.

[3] Broder, A., Kirsch, A., Kumar, R., Mitzenmacher, M., Upfal, E., and Vassilvitskii, S. The Hiring Problem and Lake Wobegon Strategies. In *ACM-SIAM SODA 08*.

[4] Burgess, J., Gallagher, B., Jensen, D., and Levine, B. N. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE INFOCOM* (April 2006).

[5] Cai, H., and Eun, D. Y. Crossing over the bounded domain: from Exponential to Power-law inter-meeting time in MANET. In *ACM MobiCom '07*.

[6] Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R., and Scott, J. Pocket Switched Networks: Real-world mobility and its consequences for Opportunistic Forwarding. Tech. Rep. UCAM-CL-TR-617, University of Cambridge, 2005.

[7] Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R., and Scott, J. Impact of Human Mobility on Opportunistic Forwarding Algorithms. *IEEE Trans. on Mobile Computing 6*, 6 (2007), 606–620.

[8] Daly, E. M., and Haahr, M. Social network analysis for routing in disconnected delay-tolerant MANETs. In *ACM MobiHoc '07* (2007), pp. 32–40.

[9] Dubois-Ferriere, H., Grossglauser, M., and Vetterli, M. Age matters: efficient route discovery in mobile ad hoc networks using encounter ages. In *ACM MobiHoc '03*, pp. 257–266.

[10] Eagle, N., and Pentland, A. Reality mining: Sensing complex social systems. *Journal of Personal and Ubiquitous Computing 10*, 4 (2005), 255–268.

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

[12] Garetto, M., Giaccone, P., and Leonardi, E. Capacity scaling in delay tolerant networks with heterogeneous mobile nodes. In *ACM MobiHoc '07* (2007), pp. 41–50.

[13] Jain, S., Fall, K., and Patra, R. Routing in a Delay Tolerant Network. In *ACM SIGCOMM '04*, pp. 145–158.

[14] Jones, E. P. C., Li, L., and Ward, P. A. S. Practical routing in delay-tolerant networks. In *WDTN '05* (2005), pp. 237–243.

[15] Karagiannis, T., Boudec, J.-Y. L., and Vojnovic, M. Power law and Exponential decay of inter contact times between mobile devices. In *MobiCom '07* (2007), ACM, pp. 183–194.

[16] Leguay, J., Friedman, T., and Conan, V. DTN routing in a Mobility Pattern Space. In *ACM WDTN '05*, pp. 276–283.

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

[18] Liu, C., and Wu, J. Scalable routing in delay tolerant networks. In *MobiHoc '07* (2007), ACM, pp. 51–60.

[19] McNett, M., and Voelker, G. M. Access and mobility of wireless pda users. *SIGMOBILE Mob. Comput. Commun. Rev. 7*, 4 (2003), 55–57.

[20] Sarafijanovic-Djukic, N., Pidrkowski, M., and Grossglauser, M. Island hopping: Efficient mobility assisted forwarding in partitioned networks. In *IEEE SECON '06*.

[21] Shiryaev, A. N. *Optimal Stopping Rules*. Springer, 2008.

[22] Spyropoulos, T., Psounis, K., and Raghavendra, C. Efficient routing in intermittently connected mobile networks: The multi copy case. *IEEE/ACM Trans. Netw. 2*, 2 (2008), 477–486.

[23] Vahdat, A., and Becker, D. Epidemic Routing for Partially Connected Ad Hoc Networks. Tech. Rep. CS-200006, Duke University, 2000.