Optimal Delay-Power Tradeoff in Sparse Delay Tolerant Networks: a preliminary study.

Giovanni Neglia
Università degli Studi di Palermo
Palermo, Italy
giovanni.neglia@tti.unipa.it

Xiaolan Zhang
University of Massachusetts
Amherst, MA, USA
ellenz@cs.umass.edu

ABSTRACT

In this paper we present a first attempt to study analytically the tradeoff between delivery delay and resource consumption for epidemic routing in Delay Tolerant Networks. We assume that the nodes cooperate in order to minimize a common cost equal to a weighted sum of the packet delivery delay and the total number of copies, which is strongly related to the power consumption. In this framework we determine the best policy each node should deploy in a very simple scenario where all the nodes have perfect knowledge of the system status. The result is used as an ideal reference to evaluate the performance of some heuristics proposed, investigating potential performance improvements and configuration criteria.

Categories and Subject Descriptors

C.2.2 [**Network Protocols**]: Routing Protocols; H.1.0 [**Models and Principles**]: Resource Tradeoffs

General Terms

Algorithms, Design, Performance

Keywords

Delay Tolerant Networks, Epidemic Routing, Performance tradeoff

1. INTRODUCTION

Epidemic routing [10] has been proposed as an approach for routing in sparse and/or highly mobile networks in which there may not be a contemporaneous path from source to destination. Epidemic routing adopts a so-called "store-carry-forward" paradigm – a node receiving a packet buffers and carries that packet as it moves, and passes the packet on to new nodes that it encounters. Analogous to the spread of infectious diseases, each time a packet-carrying node encounters a new node that does not have a copy of that packet, the carrier is said to *infect* this new node by passing on a packet copy; newly infected nodes, in turn, behave similarly. The destination receives the packet when it first meets an infected node. Epidemic routing is able to achieve minimum delivery delay at the expense of increased use of resources such as buffer space, bandwidth, and transmission power.

Variations of epidemic routing have recently been proposed in order to exploit the trade-off between delivery delay and resource consumption. For example, under probabilistic forwarding [5, 4], an infected node does not copy the packet at every encounter, but it copies the packet to other relay nodes with probability p. The higher the probability, the lower the delay and the higher the consumption of resources. When $p = 0$ the scheme reduces to direct forwarding from the source to the destination, while when $p = 1$ it reduces to standard epidemic routing. Other variations take into consideration the fact that the copies made at the start of the epidemic spreading are more useful than those made at the end. For example in K-hop schemes [6, 3] a path to the destination cannot be longer than K hops: each packet has a Time To Live (TTL), that is decreased by one from each node receiving a new copy, when the TTL is one, the node can deliver the packet only to the destination. Spray-and-wait [9, 8] (or token-based forwarding), on the other hand, limits the total number of copies that can be made for each packet using tokens. When the packet is generated, the source has T tokens available. Every time an infected node copies a packet to another node, it also transfers to the node half of its tokens. When an infected node has only one token, it can only deliver the packet to the destination. This way, the number of tokens T is the maximum number of copies in the system. Note that both K-hop and token-based forwarding behave as classic epidemic routing at the beginning of the infection, while they reduce the number of copies as the infection spreads.

Although many schemes have been proposed, nobody, to the best of our knowledge, has tried so far to study analytically the tradeoff issue. In this paper we assume that all nodes in the network cooperate (a reasonable assumption in sensor networks) in order to minimize a common cost equal to a weighted sum of the delivery delay of the packet and the total number of copies made in the system. The total number of copies made for each packet is strongly related to the transmission power consumption. In this framework we address the problem of determining the best policy each node should deploy in order to minimize the cost: the policy prescribes if the node, according to its available information, should copy or not the packet when it encounters another non-infected node. In general the policy depends on the information available to the node. In this paper we analytically determine the best policy in the most favourable case where each node has perfect knowledge of the current system state, for example it knows the current number of copies present in the system[1]. This is clearly an ideal case, as, in reality, the spreading of state information in the system is limited by the same encounter process that limits the spreading of the packet. Nevertheless it can be considered a best case reference to compare the performance of existing heuristics and estimate potential improvements. As regards this issue, we consider in this paper probabilistic forwarding and token-based forwarding: our study shows that token-based forwarding is able to achieve near-optimal performance under an appropriate setting, and our analysis of the ideal case study provides some hints to identify such setting.

In our study we mainly rely on Markovian models for the packet spreading process: the inter-meeting times between nodes are assumed to be independent and identically distributed exponential random variables. Markovian models have been used to study the performance of epidemic routing [7, 3, 4], 2-hop forwarding [3], and spray-and-wait [9, 8] and fluid models have been derived from Markovian models [12]. A support for this assumption comes from [3], where the authors consider common node mobility models (e.g., random waypoint and random direction mobility) and show that nodal inter-meeting times are nearly exponentially distributed when transmission ranges are small compared to the network's area, and node velocity is sufficiently high. This observation suggests that Markovian models of epidemic routing can lead to quite accurate performance predictions. Indeed [3] develops Markov chain models for epidemic routing and 2-hop forwarding, and derives the average source-to-destination delivery delay and the number of extant copies of a packet at the time of delivery; model predictions are validated through simulations.

2. THE MODEL

We consider a set of $N+1$ nodes with a finite transmission range moving in a closed area. We say that two nodes "meet" when they come within transmission range of each other, at which point they can exchange packets. There can be multiple source-destination pairs, but we assume that at a given time there is a single packet, eventually with many copies, spreading in the network[2]. The source of the packet can be viewed as the first carrier of a new disease, the first *infected* node. Every time it meets another node, it can decide to copy the packet to (infects) the other or not. The transmission of the packet copy requires energy, hence it reduces the lifetime of the node battery, at the same time it decreases the expected delivery delay. The new infected nodes act in the same way. As a result, the population of *susceptible* nodes –i.e., nodes without a copy of the packet– decreases over time. Once a node carrying a copy of the packet meets the destination, it passes the packet on to the destination, deletes the packet from its own buffer, and retains "packet-delivered" information (an anti-packet) which will prevent it from receiving another copy of this packet in the future; such a node has *recovered* from the disease. Different recovery schemes can be deployed [12]. In this work we ignore the recovery process because we assume that nodes know instantaneously when the packet has been delivered to the destination[3].

As we said in the introduction, we consider a Poisson process for the meetings among the nodes: the pairwise inter-meeting times are exponential random variable with rate β. [3] showed that the pairwise meeting time between nodes is nearly exponentially distributed, if nodes move in a limited region (of area A) according to common mobility models (such as the random waypoint or random direction model [2]) and if their transmission range (d) is small compared to A, and their speed is sufficiently high. The authors also derived the following formula for estimating the pairwise meeting rate β:

$$\beta \approx \frac{2wd\mathrm{E}\{V^*\}}{A}, \tag{1}$$

where w is a constant specific to the mobility model, and $\mathrm{E}\{V^*\}$ is the average relative speed between two nodes.

Let us define $n_I(t)$ the number of infected nodes at time t. The superposition of independent Poisson processes is a Poisson process with rate equal to the sum of the rates. Hence the meeting process between infected nodes and the destination is a nonhomogeneous Poisson process with rate $\beta n_I(t)$. Similarly it can be shown that the meeting process between infected nodes and susceptible nodes is a nonhomogeneous Poisson process with rate $\beta n_I(t)(N - n_I(t))$.

We want to determine the optimal policy of each node, i.e., the decision criterium according to which an infected node decides if copying or not the packet when it meets a susceptible node. The optimization goal is to minimize the following cost:

$$J = \mathrm{E}\{T_d + \gamma * M_C\}, \tag{2}$$

where T_d is the delivery delay to the destination, M_C is the total number of copies done in the network and γ is a parameter which allows one to relate time and energy consumption. The parameter γ is a design choice, the higher its value the more we give importance to the energy issue in comparison to timely delivery.

For the sake of clarity we start introducing our notation in the following simple scenario.

2.1 Simple Epidemic Spreading

Let us consider the following scenario: each infected node copies the packet every time it can, i.e., every time it meets a susceptible node, and the whole process instantaneously stops when the destination receives a copy –i.e., all the nodes are immediately informed that the destination has received the packet.

We denote the beginning of the infection (when the packet is generated at the source) as $t_0 = 0$ and consider $n_I(0) = 1$. The state can assume values $0, 1, 2, \dots N$, where $n_I(t) = 0$

[1]The rigorous definition of the system state is in Section 2.

[2]Nothing changes if we consider many packets, but we assume that the bandwidth and the buffer are large enough to assure that the different propagation processes are independent.

[3]Note that in some cases the cost of the recovery process does not affect the determination of the optimal policy. For example under Vaccine scheme [12], the anti-packet is propagated to all the nodes, hence the recovery cost is constant (equal to the cost of transmitting N anti-packets).

corresponds to the final absorbing state when the packet has been delivered to the destination. The state changes whenever an infected node meets a susceptible one or the destination. If we denote the time of the k-th meeting as t_k, then the new state is $n(t_k) = n(t_{k-1}) + 1$ for a meeting with a susceptible node, $n(t_k) = 0$ for a meeting with the destination[4].

We will show that the optimal policy depends on the transition probabilities and the average transition times. Let us define $m_{i,j}$ the probability of transition from state i to state j, and $G(i)$ the average transition time from state i. In this simple case if the state at time t_k is $n(t_k) = i$, then the probability to move to the state $i+1$ is the probability that infected nodes meet another susceptible node before the destination, i.e., $m_{i,i+1} = (N-i)/(N-i+1)$. Similarly the probability to move to state 0 is $m_{i,0} = 1/(N-i+1) = 1-m_{i,i+1}$. All the other transition probabilities are zero. The average transition time from state i is equal to $1/(\beta i(N-i+1))$.

2.2 The Optimal Forwarding Scheme under Perfect State Information

Here we introduce the decision process of the nodes: when an infected node meets a susceptible one it can decide to make a copy or not. The decision at time t_k is denoted u_k, where $u_k = c$ if the packet is copied, while $u_k = \overline{c}$ otherwise. The transition probabilities and the average transition times are now a function of the decision at time t_k as we are going to detail.

We assume perfect state information at each node, i.e., all the nodes know exactly the number of infected nodes in the system at time t ($n_I(t)$) and if the destination has already received the packet. The performance in this ideal case represents a lower bound for the cost of each real system. In the more general case we deal with a distributed stochastic optimal control problem, but under the previous assumption all the nodes have the same information, hence the distributed nature of the decision process is lost and we can consider a single controller which decides, whereas the actuator of the decision is the specific infected node. For this reason the problem can be studied as a stochastic shortest path with finite state and exponential transition time.

In order to present the derivation coherently with the common description of stochastic shortest path (see for example [1]), we need to define in a more elaborated way the state of the system. We want introduce the decisions as functions of the current state of the system. For this reason we define the state of the system at the meeting time t_k, $x_k = x(t_k)$, as the number of infected nodes before the decision, which in general changes such number. However the meeting with the destination has to be considered differently. In this case there is no decision to be made by the node and we consider that the state of the system immediately becomes zero ($x(t_k) = 0$). Formally, $x(t) = 0$ if at time t an infected node meets the destination, otherwise $x(t) = n_I(t^-) = \lim_{\tau \to t^-} n_I(t)$. Besides we assume $x(0) = 1$.

Now we are able to define clearly what is an optimal criterium. An admissible policy is the (infinite) set of functions $\pi = \{\mu_1, \dots, \mu_k, \dots\}$, where μ_k maps state x_k into controls u_k. The cost function of an admissible policy starting at time t_1 from state i is the limit of the cost from t_1 to t_K, when K diverges, i.e.:

$$
\begin{aligned}
J_\pi(i) &= \lim_{K\to\infty} \sum_{k=1}^{K-1} \mathrm{E}\left\{ \left(\hat{g}(x_k, \mu_k(x_k)) + \right.\right. \\
&\qquad \left.\left. + \int_{t_k}^{t_{k+1}} g(x_k, \mu_k(x_k)) dt \right) \Big| \, x_1 = i \right\}, \qquad (3)
\end{aligned}
$$

where $\hat{g}(x_k, \mu_k(x_k))$ is the finite cost of the decision taken at time t_k, while $g(x_k, \mu_k(x_k))$ is the cost per unit time. If we want that the total cost is expressed by equation (2), we have to consider $g(x_k, \mu_k(x_k)) = 1$, $\hat{g}(x_k, c) = \gamma$ and $\hat{g}(x_k, \overline{c}) = 0$. The total cost J_π includes also the time cost from the begin of the infection $t_0 = 0$ to the first meeting time t_1 and the cost of the last copy to the destination. It holds:

$$
\begin{aligned}
J_\pi &= \mathrm{E}\left\{ \int_{t_0}^{t_1} 1 dt \Big| \, x_0 = 1 \right\} + J_\pi(1) + \gamma = \\
&= \mathrm{E}\left\{ t_1 \Big| \, x_0 = 1 \right\} + J_\pi(1) + \gamma. \qquad (4)
\end{aligned}
$$

Note that in $[t_0, t_1)$ there is no decision from the node, hence this additional term does not depend on the policy, and we can simply consider equation (3) to identify the best policy. Similarly as regards the cost of the last copy which is constant and has to be borne in any case.

Equation (3) can be written in the form:

$$
\begin{aligned}
J_\pi(i) &= \hat{g}(i, \mu_1(i)) + G(i, \mu_1(i)) + \sum_{j=0}^{N} m_{i,j}(\mu_1(i)) J_{\pi_2}(j) \\
&= \hat{g}(i, \mu_1(i)) + G(i, \mu_1(i)) + \\
&\quad + \sum_{j=1}^{N} m_{i,j}(\mu_1(i)) J_{\pi_2}(j), \qquad (5)
\end{aligned}
$$

where J_{π_2} is the cost-to-go of the policy $\pi_2 = \{\mu_2, \mu_3, \cdots\}$ that is used from the second meeting time, $m_{i,j}(u)$ is the probability of transition from state i to state j under the decision u, and $G(i,u)$ is the average transition time from state i to another state, when the decision is u. The second equality is a consequence of zero cost in the final absorbing state ($J_\pi(0) = 0$). The following expressions for $G(i,u)$ and $m_{i,j}$ (with $i, j \neq 0$) are derived in Appendix A.

$$
G(i,u) = \begin{cases} \dfrac{1}{\beta i(N-i+1)} & \text{if } u = \overline{c}, \\ \dfrac{1}{\beta(i+1)(N-i)} & \text{if } u = c, \end{cases}
$$

$$
m_{i,j}(u) = \begin{cases} \dfrac{N-i}{N-i+1} & \text{if } j = i \text{ and } u = \overline{c}, \\ \dfrac{N-i-1}{N-i} & \text{if } j = i+1 \text{ and } u = c, \\ 0 & \text{otherwise.} \end{cases}
$$

The optimal cost function J^* is the unique solution of Bellman's equation [1]:

$$
J^*(i) = \min_{u \in \{c, \overline{c}\}} \left\{ \hat{g}(i,u) + G(i,u) + \sum_{j=1}^{N} m_{i,j}(u) J^*(j) \right\} \quad (6)
$$

A stationary policy is an admissible policy of the form $\pi = \{\mu, \mu \dots\}$. For brevity we refer to $\{\mu, \mu \dots\}$ as the sta-

[4]The packet will not be copied anymore after the delivery to the destination, and we can assume that all the copies will be deleted at each node.

tionary policy μ. A stationary policy is optimal if and only if for every state i, $\mu(i)$ attains the minimum in Bellman's equation [1].

Now we are ready to introduce the main analytical contribution of the paper.

PROPOSITION 1. *There is an optimal stationary policy μ with a threshold behavior, i.e., $\exists\, h \in \{1, 2, \dots, N\}$, such that*

$$\mu(i) = \begin{cases} c & \text{if } i < h, \\ \overline{c} & \text{otherwise,} \end{cases}$$

and the threshold h is the minimum between N and the solution of the following inequalities:

$$h(h-1) < \frac{1}{\beta\gamma} \leq h(h+1).$$

The proposition states that the optimal policy for each node is to copy the packet every time it is possible until the total number of infected nodes is equal to h ($h-1$ copies are done), or the packet is delivered to the destination.

It is interesting to note that, as far as there are enough nodes, the threshold value does not depend on the total number of nodes in the network, but simply on their pairwise meeting rate and on weight γ in the cost function. This appears evident if one considers the following problem: how many infected nodes we should start the infection with, assuming that we want to minimize the cost (2) and we have to pay γ for each infected node but the first one? The optimal number is equal to h. The total number of nodes N can only act as a constraint, when we would like to pay for more infected nodes but we do not have enough nodes, otherwise the choice of initial infected nodes does not depend on N, because these nodes will interact only with the destination and the cost depends only on the pairwise meeting rate. While the optimal policy for our original problem does not depend on N, the expected cost does. In fact the total number of nodes has an effect on the speed of the infection before reaching h infected nodes, and also on the expected number of copies which will be done before the delivery to the destination. The proof of Proposition 1 is in Appendix B.

The following proposition presents an interesting interpretation of the optimal policy.

PROPOSITION 2. *An equivalent representation for the optimal stationary policy is*

$$\mu(i) = \begin{cases} c & \text{if } \gamma + \dfrac{1}{\beta(i+1)} < \dfrac{1}{\beta i}, \\ \overline{c} & \text{otherwise.} \end{cases}$$

Hence we can think that each node compares two costs in order to decide if copying or not the packet: $1/(\beta i)$ is the expected residual (i.e., from current state) cost if the packet is not copied, while $\gamma + 1/(\beta(i+1))$ is the expected residual cost if the packet is copied only one more time. The nodes decides to copy or not respectively if the second cost or the first one is lower. While the first cost looks intuitive, we could expect the node to consider the possibility to add more than one copy (it has now the possibility to add one more, but other nodes or it itself could make more copies in the future). This particular structure of the policy comes out from the specific marginal costs of the problem. Figure 1 shows the expected time cost $\mathrm{E}\{T_d\}$ versus the number of infected nodes i ($\mathrm{E}\{T_d\} = 1/(\beta i)$), it appears that the marginal benefit from increasing the number of infected nodes by one

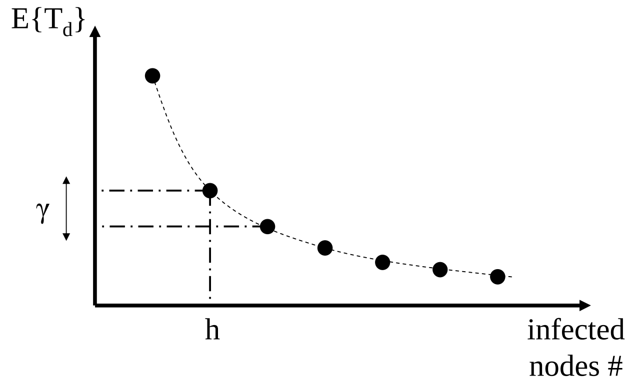

Figure 1: Expected Time Cost vs Number of Infected Nodes

Table 1: Delivery Delay and Number of Copies

Mechanism	$\mathrm{E}\{T_d\}$	$\mathrm{E}\{M_{C,del}\}$	$\mathrm{E}\{M_C\}$
opt $\gamma = 0.5$	17.62	19.42	19.42
opt $\gamma = 4$	33.34	7.64	7.64
prob 0.1%	209.18	1.09	1.09
prob 1%	176.98	1.96	1.97
prob 2%	153.53	3.02	3.14
prob 4%	102.73	5.01	5.53
prob 10%	57.61	9.90	12.51
prob 20%	39.01	17.07	25.46
prob 40%	24.53	30.88	52.98
prob 60%	19.45	39.84	71.36
prob 80%	15.22	43.33	79.08
prob 100%	12.77	48.85	87.35
token 2	114.30	1.95	1.95
token 5	57.78	4.86	4.91
token 10	29.33	8.81	9.44
token 20	18.44	15.91	18.31
token 30	17.43	21.79	27.02
token 40	16.75	27.50	34.98
token 50	15.04	29.68	41.69
token 60	16.04	33.96	48.68
token 80	14.78	37.76	57.44
token 100	14.11	40.02	64.30

decreases as i increases, while the marginal power cost γ is constant. When $i = h$ the marginal power cost exceeds the marginal benefit coming from time cost decrease. The proof of Proposition 2 is in Appendix C.

3. COMPARISON WITH HEURISTICS

In this section we compare the performance of the ideal algorithm and two heuristics we described in the introduction: probabilistic forwarding [5, 4] and Spray-and-Wait [9, 8]. Both schemes enable the trade-off between delivery delay and number of copies through a tunable parameter: the forwarding probability p and the number of token T. In this paper we consider a variant of Spray-and-Wait where two infected nodes equilibrate their number of tokens when they meet [11]. We refer to this variant as token-based forwarding.

As reference we consider $N + 1 = 101$ nodes and exponential pairwise meeting times with rate $\beta = 0.004$[5]. Also

[5]Note that, under exponential assumption, the performance

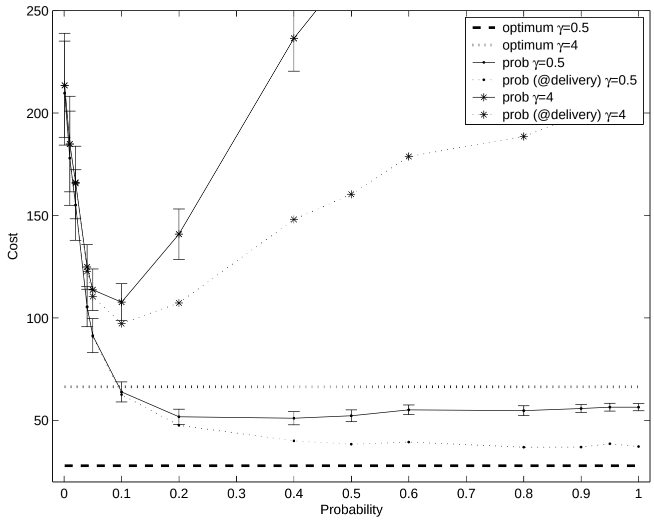

Figure 2: Probabilistic forwarding vs Optimal forwarding

when we recur to simulations we directly consider exponential inter-meeting times according to [3]. As regards the cost ($J = \mathrm{E}\{T_d + \gamma * M_C\}$) at the moment we consider $\gamma_1 = 0.5$ and $\gamma_2 = 4$. Clearly in the second case the energy consumption has a higher relevance. The corresponding thresholds for the optimal scheme can be determined from the inequalities in Proposition 1, they are $h_1 = 22$ and $h_2 = 8$ respectively. The corresponding expected costs (evaluated recursively from equations (5) and (4)) are 27.94 and 66.42. Table 1 shows also the average delivery delay $\mathrm{E}\{T_d\}$ and the average number of copies $\mathrm{E}\{M_C\}$. $\mathrm{E}\{M_{C,del}\}$ is the average number of copies at the moment of the delivery to the destination, it clearly coincides with $\mathrm{E}\{M_C\}$ in this ideal case. These numerical values have been obtained by simulations of the corresponding Markov processes, the relative width of the 95% confidence interval is below 2%.

Table 1 shows the same metrics for the two heuristics schemes for 10 different settings of p and T. Both schemes can exploit the trade-off between delivery delay and number of copies: for probabilistic (token-based) forwarding the average delivery delay decreases as $p(T)$ increases, but at the same time the number of copies increases. These values have been determined using an event-driven simulator, the relative width of the 95% confidence intervals is lower than 15%.

Now we consider the performance of the different schemes in terms of the weighted cost $J = \mathrm{E}\{T_d + \gamma * M_C\}$. In Figure 2 the two horizontal dashed lines correspond to the costs achieved by the optimal algorithm for $\gamma_1 = 0.5$ and $\gamma_2 = 4$. The solid curves are obtained from the values in Table 1 (considering also other settings), they represent the average costs that probabilistic forwarding achieves for different values of p: one curve corresponds to γ_1, the other to γ_2. For small values of p the main contribution to the cost is the delivery delay and the cost decreases as p increases, because the delivery delay decreases. On the contrary for higher

depends on the product $\beta\gamma$, hence it is not restrictive to consider a single value for β as far as we consider different values for γ.

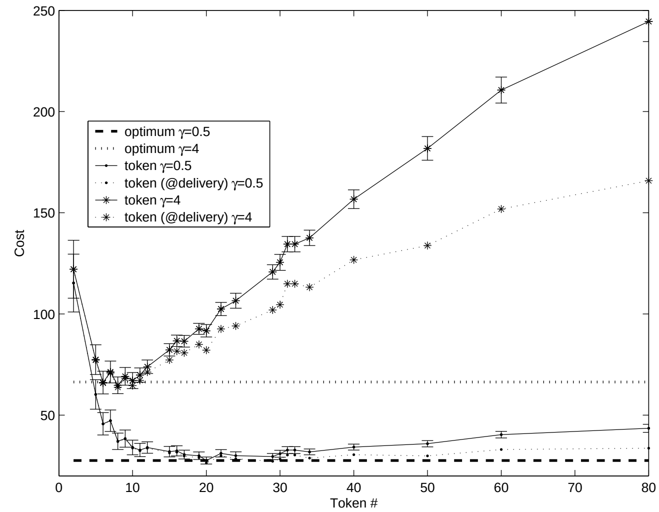

Figure 3: Token-based forwarding vs Optimal forwarding

Table 2: Configuration of Token-based forwarding on the basis of the Optimal forwarding scheme

γ	threshold (h)	Loss% ($@T = h$)	Loss% (Min,Avg)
0.25	32	15%	(4.1%,55%)
0.5	22	12%	(6.0%,41%)
1	16	11%	(5.0%,37%)
2	11	0.0%	(-0.6%,47%)
4	8	-2.6%	(-2.6%,72%)
8	6	-3.2%	(-3.2%,120%)

probability p the number of copies becomes the larger term in equation (2) and the cost increases as p increases. The slope of the curve for high p is clearly less steep for $\gamma = 0.5$ than for $\gamma = 4$. The best settings for probabilistic forwarding are $p \approx 0.4$ and $p \approx 0.1$ respectively for $\gamma = 0.5$ and $\gamma = 4$. The increase in cost in comparison to the optimal scheme is respectively 50% and 60%. The dashed curves represent the costs of probabilistic forwarding if we assume that when the packet is delivered to the destination, this information is instantaneously propagated to all the nodes which stop copying the packet. In other words, these are the costs accumulated up to the delivery time and they have been evaluated considering the average number of copies at the moment of the delivery ($\mathrm{E}\{M_{C,del}\}$) in Table 1. Under this assumption the performance loss of the scheme is reduced but it is still significant for large γ.

Figure 3 shows the corresponding curves for the token-based mechanism. It appears that this mechanism not only achieves better performance than probabilistic forwarding, but also achieves a cost which is close to optimal with a careful setting. The performance loss is almost not appreciable: the confidence intervals corresponding to the best settings of token-based forwarding include the minimum cost achievable from the optimal scheme. It is also interesting to observe that the best setting for token-based forwarding (as derived from simulations) seems strongly related to the threshold of the optimal algorithm. This is not surprising since both the threshold and the number of tokens correspond to the maximum number of infected nodes in the network. In particular, Table 2 shows the optimal threshold values (second column) for different γ and the performance loss of token-based forwarding (third column) when the number of tokens is set equal to the thresholds, i.e., $T = h$. The performance loss is evaluated as the difference between the cost under token-based forwarding and the optimal cost divided by the optimal cost[6]. The results of such configurations are compared with the minimum and the average loss over all the possible token configurations (forth column). When the cost sensitivity to γ is higher[7] (e.g. for $\gamma = 4, \gamma = 8$), such configuration criterium seems to achieves the best performance. For smaller values of γ (flat curves, like that for $\gamma = 0.5$ in Figure 3) it becomes harder to identify correctly the best setting, anyway the proposed criterium achieves good results in comparison to a blind setting, which would produce the average loss in Table 2.

4. CONCLUSIONS

This work presents the first attempt to study analytically the tradeoff between delivery delay and resource consumption for epidemic routing in Delay Tolerant Networks. We considered the ideal scenario where all the nodes have perfect knowledge of the system status and we have been able to completely characterize the optimal policy under this assumption. The optimal policy confirms the intuition that copies of the packet are more useful at the beginning of the infection.

Two heuristics have been compared with the optimal forwarding scheme: probabilistic forwarding and token-based forwarding. The preliminary results suggest that, in a realistic scenario, it is probably hard to design new schemes able to achieve better performance than token-based forwarding does; it could be more profitable to identify configuration criteria for token-based forwarding. Our study of the optimal scheme seems to provide also some directions to investigate such configuration criteria.

5. ACKNOWLEDGMENTS

This research was supported in part by Italian MIUR project Famous, National Science Foundation under award number ANI-0085848, EIA-0080119. The author Giovanni Neglia thanks Dario Bauso and prof. Raffaele Pesenti for the useful conversations about optimal stochastic control.

6. REFERENCES

[1] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, second edition, 2000.

[2] C. Bettstetter. Mobility modeling in wireless networks: categorization, smooth movement, and border effects. In *ACM SIGMOBILE Mobile Computing and Communications Review*, volume 5, Issue 3, July, 2001.

[3] R. Groenevelt, P. Nain, and G. Koole. The message delay in mobile ad hoc networks. In *Performance*, October 2005.

[6] Negative performance losses are due to simulation randomness, zero is always in the confidence intervals for such cases.

[7] This corresponds to marked minima in the curves in Figure 3 .

[4] Z. J. Haas and T. Small. A new networking model for biological applications of ad hoc sensor networks. *IEEE/ACM Transactions on Networking*, 14(1):27–40, February 2006.

[5] A. Lindgren, A. Doria, and O. Schelen. Probabilistic routing in intermittently connected networks. In *Proceedings of MobiHoc (Poster)*, 2003.

[6] G. Sharma and R. R. Mazumdar. Delay and capacity tradeoffs for wireless ad hoc networks with random mobility. Submitted for publication.

[7] T. Small and Z. J. Haas. The shared wireless infostation model - a new ad hoc networking paradigm. In *Mobihoc*, 2003.

[8] T. Small and Z. J. Haas. Resource and performance tradeoffs in delay-tolerant wireless networks. In *ACM workshop on Delay Tolerant Networking*, 2005.

[9] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In *ACM workshop on Delay-tolerant networking*, 2005.

[10] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical report, Duke University.

[11] X. Zhang, G. Neglia, J. Kurose, and D. Towsley. On the benefits of random linear coding for unicast applications in disruption tolerant networks. In *IEEE NetCod 2006*.

[12] X. Zhang, G. Neglia, J. Kurose, and D. Towsley. Performance modeling of epidemic routing. In *Networking 2006*.

APPENDIX

A. $G(i,u)$ AND $m_{i,j}(u)$

$G(i,u)$ is the average transition time from state i to another state when the decision is u. The state changes when an infected node meets a susceptible node or the destination. If $u=\overline{c}$, then the number of infected nodes does not change. The transition rate is $\beta i(N-i+1)$ and the average transition time:

$$G(i,\overline{c})=\frac{1}{\beta i(N-i+1)}.$$

Otherwise, if $u=c$ the number of infected node increases to $i+1$, the transition rate is $\beta(i+1)(N-i)$ and the transition time:

$$G(i,c)=\frac{1}{\beta(i+1)(N-i)}.$$

$m_{i,j}(u)$ is the probability of transition from state i to state j under the decision u. We consider $i,j\neq 0$ because we assume zero cost in the final absorbing state 0, hence transitions to this state do not appear in equation (5) and we do not need to evaluate transition probabilities $m_{i,0}$. If $u=\overline{c}$, the number of infected nodes keeps constant equal to i, and the state at the next meeting time is still i with probability $(N-i)/(N-i+1)$, i.e., the probability that infected node meet another susceptible node before the destination. If $u=c$, the number of infected nodes increases to $i+1$, and the state at the next meeting time is $i+1$ with probability $(N-i-1)/(N-i)$.

B. PROOF OF PROPOSITION 1

PROOF. The proof is divided into four parts:

1) we show that we can restrict our attention to stationary policy;

2) we evaluate Bellman's equation (6) and derive some useful relations;

3) we prove that the optimal stationary policy is a threshold policy;

4) we derive the threshold value.

1) Due to the markovian assumption, the decision process can take into account only the current number of infected nodes, without considering the number of previous encounters. For this reason we can restrict our attention to stationary policy μ. A stationary policy is optimal if and only if it achieves the minimum in Bellman's equation (6).

2) We evaluate Bellman's equation (6):

$$\begin{aligned}J^*(i) &= \min\Big\{\hat{g}(i,c)+G(i,c)+\sum_{j=1}^{N}m_{i,j}(c)J^*(j),\\ &\qquad \hat{g}(i,\overline{c})+G(i,\overline{c})+\sum_{j=1}^{N}m_{i,j}(\overline{c})J^*(j)\Big\}=\\ &= \min\{\gamma+G(i,c)+m_{i,i+1}J^*(i+1),\\ &\qquad G(i,\overline{c})+m_{i,i}J^*(i)\}. \qquad (7)\end{aligned}$$

We observe that under an optimal scheme if the best decision in state l is "not-copy" ($\mu(l)=\overline{c}$), then we can have at most l copies in the system, in fact the system will never reach state $l+1$ and the expected cost from state l is equal to the expected time to meet the destination from state l, i.e., $1/(\beta l)$. This can be derived from equation (7), in fact $\mu(l)=\overline{c}$ implies that the minimum is achieved with the second term, i.e.:

$$J^*(l)=G(l,\overline{c})+m_{l,l}J^*(l),$$

and after some calculations we derive

$$J^*(l)=\frac{1}{\beta l}.$$

We define $J_{\overline{c}}(i)=1/(\beta i)$ and $J_c(i)=\gamma+G(i,c)+m_{i,i+1}J^*(i+1)$. The Bellman's equation can be simply rewritten as:

$$J^*(i)=\min\{J_c(i),J_{\overline{c}}(i)\}$$

The following relation holds:

$$J_{\overline{c}}(i)=G(i,\overline{c})+m_{i,i}J_{\overline{c}}(i).$$

In fact it can be interpreted in the following way considering the regeneration process: the average time to meet the destination is equal to the average time to meet the destination or a susceptible node plus the expected time to meet the destination times the probability that the node has not met the destination ($m_{i,i}$). Similarly:

$$J_{\overline{c}}(i+1)=G(i,c)+m_{i,i+1}J_{\overline{c}}(i+1). \qquad (8)$$

3) We want to prove that an optimal policy μ is a threshold policy, more specifically there exists $0<h\leq N$ such that

$$\mu(i)=\begin{cases}c & \text{if } i<h,\\ \overline{c} & \text{otherwise}\end{cases}$$

In particular we are going to show that if

a) $J^*(i+1)=J_{\overline{c}}(i+1)$ (it is always satisfied for $i+1=N$) and

b) $J^*(i)=J_c(i)$,

then $J^*(i-k)=J_c(i-k)$ for all $k\geq 1$.

$$J^*(i)\underset{b}{=}J_c(i)=\gamma+J^*(i+1)\underset{a}{=}\gamma+J_{\overline{c}}(i+1)=\gamma+\frac{1}{\beta(i+1)},$$

and from

$$J^*(i)\underset{b}{<}J_{\overline{c}}(i),$$

it follows:

$$\gamma+\frac{1}{\beta(i+1)}<\frac{1}{\beta i}. \qquad (9)$$

Note that

$$\gamma+\frac{1}{\beta(j+1)}<\frac{1}{\beta j},\ \text{for all } j\leq i. \qquad (10)$$

For $k\geq 1$:

$$\begin{aligned}J_c(i-k) &= \gamma+G(i-k,c)+m_{i,i+1}J^*(i-k+1)\\ &\leq \gamma+G(i-k,c)+m_{i,i+1}J_{\overline{c}}(i-k+1)\\ &= \gamma+J_{\overline{c}}(i-k+1)\\ &= \gamma+\frac{1}{\beta(i-k+1)}\\ &< \frac{1}{\beta(i-k)}\\ &= J_{\overline{c}}(i-k),\end{aligned}$$

where the first inequality is a consequence of $J^*(i-k) = \min\{J_c(i-k), J_{\overline{c}}(i-k)\}$, the second equality follows from equation (8) and the second inequality follows from inequality (10).

4) Once that we proved that the optimum policy is a threshold policy, we can determine the threshold value h. It is equal to N, if $i = N-1$ satisfies equation (9): in this case the number of infected nodes can reach its maximum value N. Otherwise h is smaller than N and it is equal to the smallest value that does not satisfies equation (9), i.e.:

$$\gamma + \frac{1}{\beta h} < \frac{1}{\beta(h-1)}$$
$$\gamma + \frac{1}{\beta(h+1)} \geq \frac{1}{\beta h}.$$

The inequalities above can also be written in the following way:

$$h(h-1) < \frac{1}{\beta\gamma} \leq h(h+1).$$

□

C. PROOF OF PROPOSITION 2

PROOF. It is easy to verify that $i < h$ if and only if $\gamma + 1/(\beta(i+1)) < 1/(\beta i)$. It follows from the inequalities in Proposition 1 and from:

- $i < h$ if and only if $i(i+1) \leq h(h-1)$;
- $\gamma + \frac{1}{\beta(i+1)} < \frac{1}{\beta i}$ if and only if $i(i+1) < \frac{1}{\beta\gamma}$.

Then we can just replace the condition $i < h$ in the optimal policy definition with the equivalent one, and we get:

$$\mu(i) = \begin{cases} c & \text{if } \gamma + \dfrac{1}{\beta(i+1)} < \dfrac{1}{\beta i}, \\ \overline{c} & \text{otherwise.} \end{cases}$$

□