Performance Modeling of Epidemic Routing

Xiaolan Zhang[1], Giovanni Neglia[2], Jim Kurose[1], Don Towsley[1]

[1] University of Massachusetts, Amherst MA 01002, USA,
{ellenz,kurose,towsley}@cs.umass.edu,
[2] Università degli Studi di Palermo,
giovanni.neglia@tti.unipa.it

Abstract. In this paper, we develop a rigorous, unified framework based on Ordinary Differential Equations (ODEs) to study epidemic routing and its variations. These ODEs can be derived as limits of Markovian models under a natural scaling as the number of nodes increases. While an analytical study of Markovian models is quite complex and numerical solution impractical for large networks, the corresponding ODE models yield closed-form expressions for several performance metrics of interest, and a numerical solution complexity that does not increase with the number of nodes. Using this ODE approach, we investigate how resources such as buffer space and power can be traded for faster delivery, illustrating the differences among the various epidemic schemes considered. Finally we consider the effect of buffer management by complementing the forwarding models with Markovian and fluid buffer models.

Keywords: Delay tolerant networks, wireless ad hoc networks, epidemic routing, performance modeling, ordinary differential equations

1 Introduction

Epidemic routing [13] 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 (i.e., a special case of Delay Tolerant Network). Epidemic routing adopts a so-called "store-carry-forward" paradigm – a node receiving a packet buffers and carries that packet as it moves, passing 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 that exploit this trade-off between delivery delay and resource consumption, including K-hop schemes [9, 3], probabilistic forwarding [8, 4], and spray-and-wait [12, 11].

Early efforts evaluating the performance of epidemic routing schemes used simulation [13, 5, 8]. More recently, Markovian models have been developed to study the performance of epidemic routing [10, 3, 4], 2-hop forwarding [3], and spray-and-wait [12, 11]. Recognizing the similarities between epidemic routing and the spread of infectious diseases, [10] used ordinary differential equation (ODE) models adapted from infectious disease-spread modeling [2] to study the source-to-destination delivery delay under the basic epidemic routing scheme, and then adopted Markovian models to study other performance metrics.

In this paper, we develop a rigorous, unified framework, based on Ordinary Differential Equations (ODE), to study epidemic routing and its variations. The starting point of our work is [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, deriving the average source-to-destination delivery delay and the number of extant copies of a packet at the time of delivery. An analytical study of such Markov chain models is quite complex for even simple epidemic models, and more complex schemes have defied analysis thus far. Moreover, numerical solution of such models becomes impractical when the number of nodes is large.

We develop ODEs as a fluid limit of Markovian models such as [3], under an appropriate scaling as the number of nodes increases. This approach allows us to then derive closed-form formulas for the performance metrics considered in [3], obtaining matching results. More importantly, we are also able to use the ODE framework to further model the so-called "recovery process" (packet deletion at infected nodes, following the successful delivery to the destination), to study more complex variants of epidemic routing, and to model the performance of epidemic routing with different buffer management schemes under buffer constraints. While different recovery processes are studied also in [10, 11] using Markov chains, model simulation is first needed to determine a number of model parameters. Many of our ODE models can be analytically solved, providing closed-form formulas for the performance metrics of interest; in cases where we resort to numerical solution, the computation complexity does not increase with the number of nodes. The drawback of our ODE models is that they are used to evaluate the moments of the various performance metrics of interest, while numerical solution of Markov chain models can provide complete distributions (e.g., for the number of packet copies in the system). Simulation results show good agreement with the predictions of our ODE models.

Through our modeling studies, we obtain insights into different epidemic routing schemes. In particular, we identify rules of thumb for configuring these schemes, we show the existence of a linear relation between total number of copies sent and the buffer occupancy under certain schemes, and we demonstrate that the relative benefit of different recovery schemes depends strongly on the specific infection process. Finally our analysis of buffer-constrained epidemic routing suggests that sizing node buffers to limit packet loss is not vital as long as appropriate buffer management schemes are used.

The remainder of this paper is structured as follows. Basic epidemic routing and our basic ODE model are described and derived in Section 2, allowing one to characterize the source-to-destination delivery delay, the number of copies made for a packet, and the average buffer occupancy. In Section 3, the model is extended for three important variations of basic epidemic routing: K-hop forwarding, probabilistic forwarding and limited-time forwarding; we use these extended models to characterize the tradeoff between delivery delay and resource (buffer, power) consumption in Section 4. In Section 5, we integrate the ODE models with Markov and fluid queue models to study the effect of finite buffers, and compare different buffer management strategies. Finally in Section 6 we summarize the paper and discuss about future work. Throughout the paper, we compare our work with related efforts, where appropriate. Due to space constraint some of the derivations are in [14].

2 Basic Epidemic Routing

In this section we develop our ODE model for basic epidemic routing [13], after briefly describing epidemic routing and the scenario we are considering. We then use the model to study three different recovery techniques for limiting the number of packet copies in the network, validating these models against simulation.

We consider a set of $N + 1$ nodes with a finite transmission range moving in a closed area and different source-destination pairs. We say that two nodes "meet" when they come within transmission range of each other, at which point they can exchange packets. Let us focus on a single packet. The analogy with disease spreading is useful in describing epidemic routing. The source of the packet can be viewed as the first carrier of a new disease, the first *infected* node, and it copies the packet to (infects) every node it meets. These 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 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. We will shortly consider more sophisticated recovery schemes.

Consider now many packets spreading at the same time in the network. We assume that when two nodes meet they can exchange an arbitrary number of packets, and each node has enough buffer to store all packets (the latter assumption is relaxed in Section 5), thus allowing different infections to be considered independently. We also assume a mechanism exists so that nodes never exchange a packet if both nodes are already carrying a copy of that packet.

2.1 ODE Models for Basic Epidemic Routing

As noted earlier, [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 [1]) 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:

$$\beta \approx \frac{2wdE[V^*]}{A}, \tag{1}$$

where w is a constant specific to the mobility models, and $E[V^*]$ is the average relative speed between two nodes. Under this approximation, [3] showed that the evolution of the number of infected nodes can be modeled as a Markov chain.

We introduce our modeling approach starting from the Markov model for simple epidemic routing. Given $n_I(t)$, the number of infected nodes at time t, the transition rate from state n_I to state n_I+1 is $r_N(n_I) = \beta n_I(N-n_I)$, where N is the total number of nodes in the network (excluding the destination). If we rewrite the rates as $r_N(n_I) = N\lambda(n_I/N)(1-n_I/N)$ and assume that $\lambda = N\beta$ is constant, we can apply Theorem 3.1 in [7] to prove that, as N increases, the fraction of infected nodes (n_I/N) converges asymptotically to the solution of the following equation[3]:

$$i'(t) = \lambda i(t)(1-i(t)), \text{ for } t \geq 0 \tag{2}$$

with initial condition $i(0) = \lim_{N\to\infty} n_I(0)/N$. The average number of infected nodes then converges to $I(t) = Ni(t)$ in the sense of footnote 3. The following equation can be derived for $I(t)$ from Eq.(2):

$$I'(t) = \beta I(N-I), \tag{3}$$

with initial condition $I(0) = Ni(0)$. Such an ODE, which we have shown results as a fluid limit of a Markov model as N increases, has been commonly used in epidemiology studies, and was first applied to broadcast in mobile ad hoc network in [6], epidemic routing in [10], as a reasonable approximation.

We remark that 1) the initial population of infected nodes must scale with N, and 2) the pairwise meeting rate scales as $1/N$. Eq.(1) also provides insight into the physical interpretation of the meeting rate scaling, in particular one can consider that the area A increases with N, keeping node density constant. In the following we will consider Eq.(3) with initial condition $I(0) = 1$, which corresponds to an initial fraction of infected nodes $i(0) = 1/N$. Despite the "small" number of initial infected nodes, we will see via our simulation results that the approximation is a good one. We also note that Eq.(3), as well as other related equations we will derive shortly, can also be obtained in a different manner from Markovian models by neglecting terms related to higher moments [2, 14].

2.2 Delay under epidemic routing

Let T_d be the packet delivery delay, i.e., the time from when a packet is generated at the source to the time when it is first delivered to the destination, and denote its Cumulative Distribution Function (CDF) by $P(t) = Pr(T_d < t)$. Under the

[3] Formally, $\forall\epsilon > 0$, $\lim_{N\to\infty} \text{Prob}\{|\sup_{s\leq t}\{n_I(s)/N - i(s)\}| > \epsilon\} = 0$

same scaling and approximations considered earlier, we can derive the following equation for $P(t)$: $P'(t) = \lambda i(1-P)$ [14], and in a similar manner

$$P'(t) = \beta I(1-P). \tag{4}$$

Eq.(4) was proposed in [10], based on an analogy with a Markov process. Solving Eq.(3) and Eq.(4) with $I(0) = 1, P(0) = 0$, we get

$$I(t) = \frac{N}{1+(N-1)e^{-\beta N t}}, \qquad P(t) = 1 - \frac{N}{N-1+e^{\beta N t}}$$

From $P(t)$, the average delivery delay can be explicitly found in closed form as $E[T_d] = \int_0^\infty (1-P(t))dt = \ln N/(\beta(N-1))$. The average number of copies of a packet in the system when the packet is delivered to the destination, $E[C_{ep}]$, can also be derived [14]: $E[C_{ep}] = \int_0^\infty I(t)P'(t)dt = \frac{N-1}{2}$.

We note that while [3] obtained the same result for the number of copies, derived the Laplace-Stieltjes Transform (LST) of the delay, and from the LST found the following asymptotic expression for the average delay as $N \to \infty$: $\frac{1}{\beta(N-1)}(\ln N + \gamma + O(\frac{1}{N}))$, the derivation is much simpler using our ODE model.

2.3 Recovery from infection

Next we study the total number of packet copies sent, and the packet's average storage requirement under the recovery schemes proposed in [4]. Clearly, once a node delivers a packet to the destination, it should delete the packet from its buffer to save storage space and prevent the node from infecting other nodes. Moreover, to avoid being reinfected by the packet, a node can store a so-called "anti-packet" once it delivers a packet to the destination. We refer to this scheme as IMMUNE scheme. A more aggressive approach towards deleting obsolete copies is to propagate anti-packets among nodes. An anti-packet can be propagated only to infected nodes (which we will term as IMMUNE_TX scheme), or to both infected and susceptible nodes (VACCINE scheme).

Similar to our earlier analysis in Section 2.1, we can derive ODEs that take into account the recovery process as the limit of Markov models [14], with the additional consideration that we need to scale the number of destinations n_D in a manner similar to the scaling of the number of initially infected nodes, i.e. $\lim_{N\to\infty} n_D/N = d$. For example, if we consider the IMMUNE scheme, the number of infected and recovered nodes should be respectively close to $I(t)$ and $R(t)$, which are solutions of the following equations:

$$I'(t) = \beta I(N-I-R) - \beta ID, \qquad R'(t) = \beta ID$$

where D is the number of destinations and we consider $I(0) = 1, R(0) = 0, D = 1$. This model allows us to evaluate the average number of times that a packet is copied during its lifetime, $E[G_{ep}]$. In fact $E[G_{ep}] = \lim_{t\to\infty} I(t) + R(t) - I(0) - R(0)$ and a good approximation can be found through the previous equations by expressing I as a function of R, without the need to solve for $I(t)$ and $R(t)$ (see Table 1 for results and [14] for a detailed derivation). Analogous ODEs can be derived for the IMMUNE_TX and VACCINE schemes, and a closed formula can be derived for $E[G_{ep}]$ for the IMMUNE_TX scheme (Table 1). Numerical solutions are needed for the VACCINE scheme.

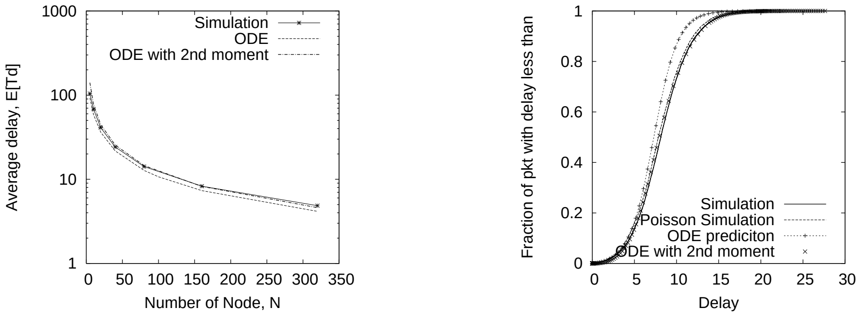

(a) Average delay for different N (b) CDF of delay with $N = 160$

Fig. 1. Delay under epidemic routing

We next consider the average storage requirement in the case of $N+1$ unicast flows, with each node being the source of one flow and destination for one other flow, and each flow generating packets with Poisson rate λ. Denote by L the average packet lifetime (the time from when the packet is generated by the source node to when all copies of the packet are removed from the system). The average number of copies of a packet in the system during its lifetime is given by $\int_0^\infty I(t)dt/L$, where $I(t)$ is the solution to the ODEs that include the recovery process. As the total arrival rate of new packerts to the system is $(N+1)\lambda$, by Little's law, the average number of packets in the system is $(N+1)\lambda L$. Therefore the average total buffer occupancy in the whole network is given by $E[Q_{total}] = (\int_0^\infty I(t)dt/L)(N+1)\lambda L = \int_0^\infty I(t)dt(N+1)\lambda$, and the per-node buffer occupancy is thus $E[Q] = \lambda \int_0^\infty I(t)dt$.

Modeling a node's buffer as an $M/M/\infty$ queue gives the same result [14] and shows a linear relationship between the average buffer occupancy and the number of copies made when IMMUNE is used. In fact, given that each packet is copied $E[G_{ep}]$ times, each flow generates relay traffic of rate $E[G_{ep}]\lambda$, and the total rate of relay traffic in the network is $E[G_{ep}]\lambda(N+1)$ (as there are $N+1$ flows). This traffic is equally divided among the $N+1$ nodes, hence the arrival rate of relay packets to each node is $E[G_{ep}]\lambda$, and the total packet arrival rate is $\lambda(1+E[G_{ep}])$. If a copy is deleted only when the node meets the destination[4], the service rate is $1/\beta$ and the average buffer occupancy is $E[Q] = \frac{\lambda}{\beta}(1+E[G_{ep}])$.

2.4 Model Validation

Throughout the paper, we validate our models using a simulator we developed to simulate the epidemic routing scheme and its variations under various mobility models. Here, we validate our models using a specific setting considered in [3]: nodes move according to random direction model within a 20×20 terrain. The transmission range of the node is chosen to be 0.1. The node speed is chosen uniformly in the range 4-10, and the mean trip duration is 1/4. The pair-wise meeting rate for this setting is found to be $\beta = 0.00435$ using the formula in [3].

[4] This is the case under IMMUNE for the basic epidemic routing, and also for two other schemes we are going to consider: probabilistic and K-hop forwarding.

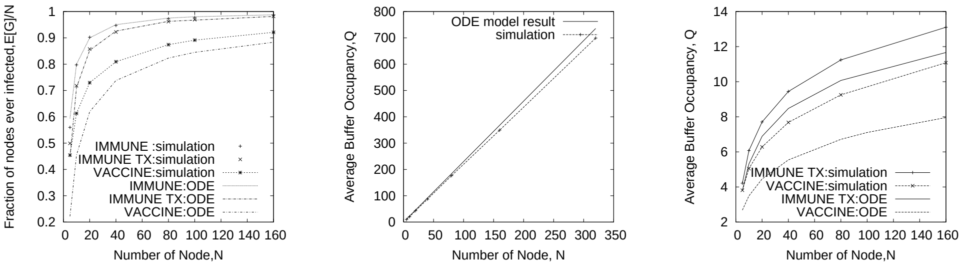

(a) $E[G_{ep}]/N$ (b) Buffer occupancy under IMMUNE (c) Buffer occupancy under IMMUNX_TX and VACCINE

Fig. 2. Copies sent and buffer occupancy under epidemic routing

We vary the number of nodes, N, and let each flow generate packets with Poisson rate $\lambda = 0.01$. The simulation is run long enough so that 100 packets are generated for each flow. The mean and CDF of the delivery delay obtained from the simulation are compared with the model results in Fig.1. We observe that the model is able to accurately predict the delivery delay, capturing the performance trend as N increases, with a slightly larger discrepancy in the CDF. To investigate modeling errors, we run another set of simulations with nodes meeting according to a Poisson process with rate $\beta = 0.00435$ (i.e., we set the meeting rate in the simulation to exactly match the model's meeting rate) and the results of the two sets of simulations are very close (Fig.1.(b)). We thus conjecture that the prediction errors are mainly due to the small number of initial infected nodes. We also use a moment-closure technique to derive a ODE system involving second moments [14]. The modified ODE provides a better prediction of average delivery delay and the CDF of delivery delay (Fig.1).

For the different recovery schemes, Fig.2 plots $E[G_{ep}(N)]/N$, and the average buffer occupancy as predicted by the model and obtained from simulation. We find that the ODE models are more accurate for IMMUNE than for VACCINE. In some sense, any error in the infection process modeling is amplified by the exponentially fast recovery of VACCINE. We observe that IMMUNE_TX only slightly reduces the number of copies sent for each packet, while VACCINE further reduces the number of copies sent. The reduction in buffer requirements is similar for IMMUNE_TX and VACCINE.

3 Extended Model

Although the recovery schemes discussed in the previous section lead to substantial differences in buffer and power requirements, they all achieve the minimum delay. The following schemes allow one to trade-off timely delivery with resource consumption.

K-Hop forwarding: Under K-hop forwarding, a packet can traverse at most K hops to reach the destination. We can model these schemes by introducing $K-1$ ODEs, describing the evolution of the number of nodes infected by i-hop paths for $1 \leq i < K$. For example in 2-hop forwarding only the source node can copy the packet to nodes other than the destination, hence the packet spreading

Table 1. Summary of closed-form expressions obtained for different schemes

Schemes	$I(t)$ $P(t)$	$E[T_d]$	C,G
Epidemic	$I(t)=\frac{N}{1+(N-1)e^{-\beta Nt}}$ $P(t)=1-\frac{N}{N-1+e^{\beta Nt}}$	$\frac{\ln N}{\beta(N-1)}$	$C=\frac{N-1}{2}$, $G\approx N-1$ (IM) $G=\frac{N-3+\sqrt{N^2-2N+5}}{2}$ (IM_TX)
2-hop	$I(t)=N-(N-1)e^{-\beta t}$ $P(t)=1-e^{N-1-\beta Nt-(N-1)e^{-\beta t}}$	$\frac{1}{\beta}\sqrt{\frac{\pi}{2}}\frac{1}{\sqrt{N-1}}$	$C=\sqrt{\frac{\pi}{2}}\sqrt{N}$, $G=\frac{N-1}{2}$
Prob. Fwding	$I(t)=\frac{N}{1+(N-1)e^{-p\beta Nt}}$ $P(t)=1-(\frac{N}{N-1+e^{p\beta Nt}})^{1/p}$	$[\frac{\ln(N)}{\beta(N-1)},\frac{\ln(N)}{\beta p(N-1)}]$	$C=\frac{p(N-1)}{1+p}$

rate equals to the rate that the source node meets susceptible nodes. This leads to the following equation for $I(t)$: $I'(t)=\beta(N-I)$ with $I(0)=1$.

Probabilistic forwarding: Under probabilistic forwarding, each relay node accepts a packet forwarded from an infected node with probability p, resulting in an effective infection rate of $p\beta$, so we have $I'(t)=\beta pI(N-I)$ with $I(0)=1$.

Limited-Time forwarding: For limited-time forwarding, when a *relay* node accepts a packet copy from an infected node, it starts a timer with an exponential random timeout value with mean $1/\mu$. When the timer expires, the node deletes the copy and stores the corresponding anti-packet so that it will not be infected by that packet again. Let $I_r(t)$ be the average number of infected *relay* nodes at time t, and $T(t)$ be the number of timed-out nodes at time t, then we have:

$$I_r'(t)=\beta(I_r+1)(N-I_r-T-1)-\mu I_r, \qquad T'(t)=\mu I_r$$

with $I_r(0)=0, T(0)=0$. A variant of this scheme is studied in [14].

With the above ODEs, the packet delivery delay is then found by $P'(t)=\beta I(1-P)$ for 2-hop and probabilistic forwarding, and $P'(t)=\beta(I_r+1)(1-P)$ for limited-time forwarding, both with $P(0)=0$. We solve the above ODEs either analytically or numerically, and then extend them to consider the recovery process. We also perform simulations to validate the models [14]. The main results are summarized in Table 1.

4 Performance Trade-off

In this section, we use our ODE models to quantitatively explore the performance trade-offs offered by the various epidemic routing schemes. Previous work [4, 11] investigated the buffer-delay trade-off by varying the number of nodes. However, we believe that the number of nodes is often given, and it is consequently more important to evaluate the performance trade-offs achieved by different schemes and/or understand how performance changes when configurable parameter values change. In terms of the power-delay trade-off, previous work [11, 12] only considered the trade-off achieved by a special scheme that enforces a fixed number of copies (and hence energy consumption). Our results are mainly based on numerical solution of ODE (for $N=100$, $\beta=0.00435$, $\lambda=0.01$), but also on the asymptotic formulas we derived.

Fig.3.(a) and Fig.3.(b) respectively plot the delay-versus-buffer-occupancy and the delay-versus-number-of-copies-sent trade-offs for the IMMUNE scheme. In the figure, the two singleton points correspond to 2-hop and 3-hop forwarding,

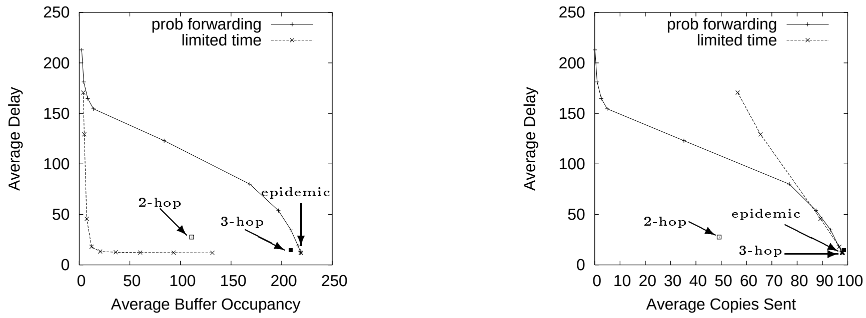

(a) delay vs buffer occupancy tradeoff (a) delay vs number of copies sent tradeoff

Fig. 3. Comparison with IMMUNE recovery

while two curves have been obtained for probabilistic forwarding and limited-time forwarding (without reinfection) respectively; for these curves, each point corresponds to a different value of the copy probability, p, and the mean timeout interval, $1/\mu$ (the values are shown in Table 2).

Table 2. Settings considered for Limited-Time and Probabilistic forwarding

Timeout ($1/\mu$)	1	2	5	10	20	40	80	160	320	
Probability (p,%)	0.1	0.5	0.8	1	2	5	10	20	50	80

Let us first consider the delay-versus-buffer-occupancy trade-off. One can reduce the buffer occupancy by decreasing p or $1/\mu$, but at the same time the delay will increase[5]. Limited-time forwarding appears to be the best choice when limiting buffer occupancy is the main concern. As a thumb rule, one can choose $1/\mu \approx 2E[T_d]$ ($= 20$ in this specific setting). This choice significantly reduces the buffer occupancy in comparison to basic epidemic routing (about one tenth), with a negligible increase in the delivery time.

Fig. 3.(b) shows that the curves for probabilistic and K-hop forwarding are similar to the delay-versus-buffer trade-off curves. This is due to the proportionality between number of copies and buffer occupancy we have shown in Section 2.3. We observe that if power saving is of primary concern, then probabilistic forwarding appears to be the best choice. For example, with $p = 0.008$ the average delay is about 30% less in comparison to the non-relaying case, while the average number of copies is only 3.5. K-hop forwarding offers intermediate performance, without sacrificing either of the two metrics.

Fig. 4 shows the same trade-offs for the VACCINE scheme. We observe that for different schemes, different performance improvements are achieved by VACCINE: in particular, the largest improvement is achieved for probabilistic forwarding, followed by K-hop forwarding, and then limited-time forwarding. The relatively small improvement for limited-time forwarding is due to its intrinsic recovery feature: nodes automatically recover as the timer expires and they cannot be reinfected. The explanation is more complex for the probabilistic and K-hop forwarding schemes. Because of the two counteracting processes – the counter-infection due to anti-packets spreading and the ongoing packet infection

[5] Intuitively these schemes behave as the original epidemic model as $p \to 1$ and $1/\mu \to \infty$, whereas $p \to 0$ and $1/\mu \to 0$ correspond to a scenario without any relay: the packet is delivered directly from the source to the destination.

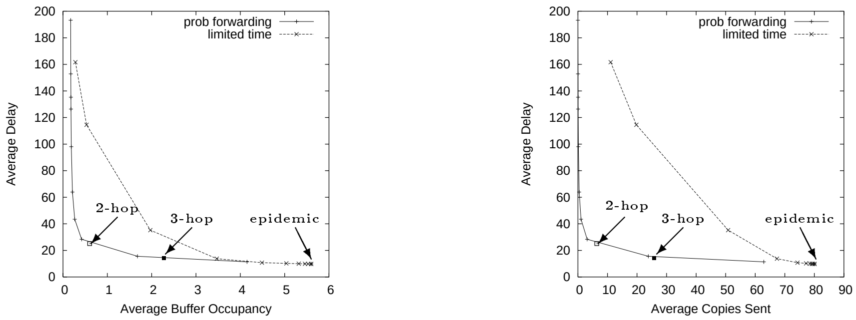

(a) delay vs buffer occupancy tradeoff (a) delay vs number of copies sent tradeoff

Fig. 4. Comparison with VACCINE recovery

– the net recovery speed depends not only on the recovery scheme but also on the specific infection process. Given the same average delivery delay, when the recovery process starts, the average number of nodes infected and the current infection rate are higher under probabilistic forwarding (its infection rate is exponential, hence in the long term it is faster than K-hop). For this reason, we expect the IMMUNE recovery process to be significantly "longer" for probabilistic routing, leading to larger buffer occupancy and more copies. Conversely under VACCINE, the recovery process is much shorter, the buffer occupancy is mainly determined by the initial infection process (before the delivery), and the difference becomes much smaller, as shown in Fig.4.

5 Epidemic Routing under Constrained Buffer

Thus far, we have assumed that each node has sufficient space to store all packets. Realistically, however, mobile nodes often have limit storage due to cost and form factor. Sizing the buffer to limit end-to-end packet losses due to buffer overflow in store-carry-forward networks is hard. For example, [4] studied buffer occupancy variability for the purpose of buffer sizing, but their model required an empirical distribution obtained from simulation. In this section, we examine the performance of epidemic routing when each node can store at most B packets. We consider three buffer management strategies: *(i) droptail* where newly arriving packets are dropped if the buffer is full (previously studied in [13] through simulation), *(ii) drophead* where the oldest packet in the buffer is dropped to accept newly arriving packets, and *(iii)* source-prioritized drophead, *drophead_sp*, which gives priority to packets arriving directly from the source. We describe the model for drophead_sp here; a full analysis can be found in [14].

Under drophead_sp, when a packet arrives to a full buffer, the node discards the oldest relay packet (i.e., a packet it has received from other node) to make space for the new packet. If all buffered packets are source packets, and the arriving packet is a source packet, the oldest source packet is deleted. Relay packets arriving to a buffer filled with source packets are not accepted. Therefore, given P_f, the probability that a node's buffer is filled with source packets, the effective infection rate is then $\beta(1 - P_f)$. P_f can be derived by modeling the number of node-buffered source packets as a Markov chain.

As before, we focus on the spreading of a single packet. Let $\overline{G}_{dhs}$ be the average number of copies made for each packet. In the source node, the copy of this packet becomes older at rate λ, the rate at which new source packets arrive. In infected relay nodes, the packet becomes older whenever another packet arrives, with rate $(\overline{G}_{dhs}+1)\lambda$ (this is the total packet arrival rate to a node by an argument similar to that in Section 2.3). Let $I_j^s(t)$ be the probability that the packet is the j-th newest source packet in the source node's buffer, $I_j(t)$ be the average number of infected *relay* nodes where the copy is the j-th newest packet in the buffer, $S(t)$ be the average number of susceptible nodes, and $D(t)$ be the average number of nodes that have dropped the packet. We can then use the following ODEs to model packet spreading:

$$\begin{array}{ll}
S'(t) = -\beta(1-P_f)S\sum_i(I_i^s+I_i) & \\
I_1'(t) = \beta(1-P_f)S\sum_i(I_i^s+I_i) - (\overline{G}_{dhs}+1)\lambda I_1 & \\
I_j'(t) = (\overline{G}_{dhs}+1)\lambda(I_{j-1}-I_j), \quad 2 \le j \le B & \\
{I_1^s}'(t) = -\lambda I_1^s, & {I_j^s}'(t) = \lambda(I_{j-1}^s - I_j^s), \quad 2 \le j \le B \\
D'(t) = (\overline{G}_{dhs}+1)\lambda I_B + \lambda I_B^s, & P'(t) = \beta\sum_i(I_i^s+I_i)(1-P)
\end{array}$$

The sums above are for $1 \le i \le B$. The initial conditions are given by: $S(0) = N-1$, $I_1^s(0) = 1$, $I_j^s(0) = 0$, for $j = 2, ...B$, $I_k(0) = 0$, for $k = 1, ..., B$, $D(0) = 0$, $P(0) = 0$. We find $\overline{G}_{dhs}$ by solving the following fixed-point problem using a binary search algorithm: given $\overline{G}_{dhs}$, we numerically solve the corresponding extended ODE model (including the recovery process) and calculate the accumulated amount of flow from state S to I_1, i.e., $\overline{G}_{dhs}$.

We have simulated these schemes, using the same setting as before ($N = 100, \lambda = 0.01$, $\beta = 0.00435$), with different buffer size $B = 5, 10, 20$, and compared our ODE results with simulation. Table 5 tabulates the loss probabilities. We observe that the models provide reasonable loss probability predictions, and accurately reflect the relative performance of the three dropping schemes. The shape of the delay distribution probability function is also well-captured by the model [14]. We observe that naive droptail performs poorly. Drophead provides fast infection, as relay packets are always accepted; however, significant packet losses are incurred for $B \le 10$. With drophead_sp, although the infection spreads slower, more packets are delivered. If the packet rate is so high that the buffer can only hold its own source packets, drophead_sp degenerates to direct source-destination transmission. Note that with infinite buffers, the average buffer occupancy for this setting is over 200 (Fig. 2.(b)). Our results here suggest that similar performance can be achieved by drophead and drophead_sp with a much smaller buffer size, equal to only 20 packets.

Table 3. Loss Probability Under Constrained Buffer

Buffer size	simulation/model	droptail	drophead	drophead_sp
5	simulation	0.9696	0.2234	0.0536
	model	0.8544	0.0928	0.0079
10	simulation	0.9471	0.0315	0.0
	model	0.7891	0.0088	0.0
20	simulation	0.899	0.0016	0.0
	model	0.7011	0.0	0.0

6 Summary and Future Work

In this paper, we proposed a unified framework based on ODEs to study the performance of epidemic routing and its variations. Using these models, we obtained a rich set of quantitative results on the delivery delay, number of copies sent, and buffer requirements (and the tradeoffs of these performance metrics) under various schemes. We further considered buffer-constrained case, and showed that with appropriate buffer management schemes, a much smaller buffer can be used with negligible effect on delivery performance. In the future, we plan to investigate schemes for deleting anti-packets and the overhead of anti-packets.

Acknowledgment

This research was supported in part by the National Science Foundation under the Engineering Research Centers Program, Award number EEC-0213747001, and Italian MIUR project Famous. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

1. 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.
2. D.J. Daley and J. Gani. *Epidemic Modelling.* Cambridge University Press, 1999.
3. R. Groenevelt, P. Nain, and G. Koole. The message delay in mobile ad hoc networks. In *Performance*, October 2005.
4. Z. J. Haas and T. Small. A new networking model for biological applications of ad hoc sensor networks. To appear in IEEE/ACM Transactions on Networking.
5. P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet. In *ASPLOS-X*, 2002.
6. A. Khelil, C. Becker, J. Tian, and K. Rothermel. An epidemic model for information diffusion in manets. In *Proceedings of MSWiM*, 2002.
7. T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump markov processes. *Journal of Applied Probabilities*, pages 49–58, 1970.
8. A. Lindgren, A. Doria, and O. Schelen. Probabilistic routing in intermittently connected networks. In *ACM Mobicom (poster session)*, 2003.
9. G. Sharma and R. R. Mazumdar. Delay and capacity tradeoffs for wireless ad hoc networks with random mobility. Submitted for publication.
10. T. Small and Z. J. Haas. The shared wireless infostation model - a new ad hoc networking paradigm. In *Mobihoc*, 2003.
11. T. Small and Z. J. Haas. Resource and performance tradeoffs in delay-tolerant wireless networks. In *ACM workshop on Delay Tolerant Networking*, 2005.
12. 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.
13. A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-200006, Duke University, April 2000.
14. X. Zhang, G. Neglia, J. Kurose, and D. Towsley. Performance modeling of epidemic routing. Technical Report 2005-44, UMASS Computer Science. ftp://gaia.cs.umass.edu/pub/Zhang05_epidemic_TR.pdf.