

Optimal forwarding in Delay Tolerant Networks with Multiple Destinations

Chandramani Singh, Eitan Altman, Anurag Kumar, Rajesh Sundaresan

► **To cite this version:**

Chandramani Singh, Eitan Altman, Anurag Kumar, Rajesh Sundaresan. Optimal forwarding in Delay Tolerant Networks with Multiple Destinations. *IEEE/ACM Transactions on Networking*, IEEE/ACM, 2013, 21 (6), pp.1812 - 1826. <10.1109/TNET.2012.2233494>. <hal-00918679>

HAL Id: hal-00918679

<https://hal.inria.fr/hal-00918679>

Submitted on 14 Dec 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Forwarding in Delay Tolerant Networks with Multiple Destinations

Chandramani Singh, *Student Member, IEEE*, Eitan Altman, *Fellow, IEEE*, Anurag Kumar, *Fellow, IEEE*, and Rajesh Sundaresan, *Senior Member, IEEE*

Abstract—We study the trade-off between delivery delay and energy consumption in a delay tolerant network in which a message (or a file) has to be delivered to each of several destinations by epidemic relaying. In addition to the destinations, there are several other nodes in the network that can assist in relaying the message. We first assume that, at every instant, all the nodes know the number of relays carrying the message and the number of destinations that have received the message. We formulate the problem as a controlled continuous time Markov chain and derive the optimal closed loop control (i.e., forwarding policy). However, in practice, the intermittent connectivity in the network implies that the nodes may not have the required perfect knowledge of the system state. To address this issue, we obtain an ODE (i.e., a deterministic fluid) approximation for the optimally controlled Markov chain. This fluid approximation also yields an asymptotically optimal open loop policy. Finally, we evaluate the performance of the deterministic policy over finite networks. Numerical results show that this policy performs close to the optimal closed loop policy.

I. INTRODUCTION

Delay tolerant networks (DTNs) [1] are sparse wireless ad hoc networks with highly mobile nodes. In these networks, the link between any two nodes is up when these are within each other’s transmission range, and is down otherwise. In particular, at any given time, it is unlikely that there is a complete route between a source and its destination.

We consider a DTN in which a short message (also referred to as a *packet*) needs to be delivered to multiple (say M) destinations. There are also N potential relays that do not themselves “want” the message but can assist in relaying it to the nodes that do. At time $t = 0$, N_0 of the relays have copies of the packet. All nodes are assumed to be mobile. In such a network, a common technique to improve packet delivery delay is *epidemic* relaying [2]. We consider a controlled relaying scheme that works as follows. Whenever a node (relay or destination) carrying the packet meets a relay that does not have a copy of the packet, then the former has the option of either copying or not copying. When a node that has the packet meets a destination that does not, the packet can be delivered.

This is an extended version of a paper that appeared in WiOpt 2011.

This work was supported by the Indo-French Centre for the Promotion of Advanced Research (IFCPAR) Project 4000-IT-1, by DAWN (an Associates program of INRIA, France), and by the Department of Science and Technology, Government of India.

Chandramani Singh, Anurag Kumar and Rajesh Sundaresan are with the Department of Electrical Communication Engineering Indian Institute of Science Bangalore, India (email: {chandra, anurag, rajesh}@ece.iisc.ernet.in). Eitan Altman is with INRIA, Sophia-Antipolis, France (email: Eitan.Altman@sophia.inria.fr).

We want to minimize the delay until a significant fraction (say α) of the destinations receive the packet; we refer to this duration as *delivery delay*. Evidently, delivery delay can be reduced if the number of carriers of the packet is increased by copying it to relays. Such copying cannot be done indiscriminately, however, as every act of copying between two nodes incurs a transmission cost. Thus, we focus on the problem of the control of packet forwarding.

A. DTNs with Multiple Destinations

DTNs are commonly envisaged for certain applications involving personal mobile devices; such applications have two characteristics: (i) they can work with infrastructure-free direct communication between the devices, and (ii) can tolerate moderate delays. Such applications typically involve spreading (or, delivering) a content to a set of nodes. The following are three examples of such applications that involve the spread of the same message to multiple destinations, i.e., intended recipients. A widely studied problem is the spread of influence in social networks, e.g., spread of a popular entertainment video clip, or diffusion of an alert about a free medical “camp” in a developing country, or an advertisement of a new product. Venkatramanan and Kumar [3] study the joint evolution of content popularity and delivery in mobile peer-to-peer (P2P) networks. The creator of the content aims to maximize its popularity at one level of the model, while at another level the parameters of the spreading process govern the dissemination of the content to those who desire it. Karnik and Dayama [4] study efficient campaigning to yield a desired level of “buzz” at a specified time, and also the diffusion of system-wide traffic updates or security alerts through wireless vehicular networks. In another work, Khouzani et al. [5] study the dissemination of security patches in mobile wireless networks, and seek optimal trade-offs between security risks and resource consumption.

While each instance has its own resource constraints and performance objectives, we consider one such multiple destination problem. The introduction of multiple sources and destinations also facilitates study of scaled network dynamics that yields an asymptotically optimal open loop policy.

B. Related Work

Analysis and control of DTNs with a single-source and a single-destination has been widely studied. Groenevelt et al. [6] modeled epidemic relaying and two-hop relaying using Markov chains. They derived the average delay and the number

of copies generated until the time of delivery. Zhang et al. [7] developed a unified framework based on ordinary differential equations (ODEs) to study epidemic routing and its variants.

Neglia and Zhang [8] were the first to study the optimal control of relaying in DTNs with a single destination and multiple relays. They assumed that all the nodes have perfect knowledge of the number of nodes carrying the packet. Their optimal closed loop control is a threshold policy - when a relay that does not have a copy of the packet is met, the packet is copied if and only if the number of relays carrying the packet is below a threshold. Due to the assumption of complete knowledge, the reported performance is a lower bound for the cost in a real system.

Spyropoulos et al. [9] introduced the *Spray and Wait* and *Spray and Focus* routing algorithms for intermittently connected mobile networks. Their algorithms disseminate the message to a predetermined number of relays, and then rely on direct delivery from any of these nodes to the destination. In this paper, we propose a finer control of spraying, that depends on the instantaneous network state. Consequently, our proposed policy is expected to outperform those in [9].

Altman et al. [10] addressed the optimal relaying problem for a class of *monotone relay strategies* which includes epidemic relaying and two-hop relaying. In particular, they derived *static* and *dynamic* relaying policies. Altman et al. [11] considered optimal discrete-time two-hop relaying. They also employed stochastic approximation to facilitate online estimation of network parameters. In another paper, Altman et al. [12] considered a scenario where active nodes in the network continuously spend energy while *beaconing*. Their paper studied the joint problem of node activation and transmission power control. These works ([10], [11], [12]) heuristically obtain fluid approximations for DTNs and study open loop controls. Li et al. [13] considered several families of open loop controls and obtain optimal controls within each family.

Deterministic fluid models expressed as ordinary differential equations have been used to approximate large Markovian systems. Kurtz [14] obtained sufficient conditions for the convergence of Markov chains to such fluid limits. Darling [15] and subsequently, Darling and Norris [16] generalized Kurtz's results. Darling [15] considered the scenario when the Markovian system satisfies the conditions in [14] only over a subset. He showed that the scaled processes converge to a fluid limit until they exit from this subset. Darling and Norris [16] generalized the conditions for convergence, e.g., uniform convergence of the mean drifts of Markov chains and Lipschitz continuity of the limiting drift function, prescribed in [14]. Gast and Gaujal [17] addressed the scenario where the limiting drift functions are not Lipschitz continuous. They proved that under mild conditions, the stochastic system converges to the solution of a differential inclusion. Gast et al. [18] studied an optimization problem on a large Markovian system. They showed that solving the limiting deterministic problem yields an asymptotically optimal policy for the original problem.

C. Our Contributions

We formulate the problem as a controlled continuous time Markov chain (CTMC) [19], and obtain the optimal pol-

icy (Section III). The optimal policy relies on complete knowledge of the network state at every node, but availability of such information is constrained by the same connectivity problem that limits packet delivery. In the incomplete information setting, the decisions of the nodes would have to depend upon their beliefs about the network state. The nodes would need to update their beliefs continuously with time, and also after each meeting with another node. Such belief updates would involve maintaining a complex information structure and are often impractical for nodes with limited memory and computation capability. Moreover, designing closed loop controls based on beliefs is a difficult task [20], even more so in our context with multiple decision makers and all of them equipped with distinct partial information.

In view of the above difficulties, we adopt the following approach. We show that when the number of nodes is large, the optimally controlled network evolution is well approximated by a deterministic dynamical system (Section IV). The existing differential equation approximation results for Markovian systems [14], [15] do not directly apply, as, in the optimally controlled Markov chain that arises in our problem, the mean drift rates are discontinuous and do not converge uniformly. We extend the results to our problem setting in our Theorem 4.1 in Section IV. Note that the differential inclusion based approach of Gast and Gaujal [17] is not directly applicable in our case, as it needs uniform convergence of the mean drift rates. The limiting deterministic dynamics then suggests a deterministic control that is asymptotically optimal for the finite network problem, i.e., the cost incurred by the deterministic control approaches the optimal cost as the network size grows. We briefly consider the analogous control of two-hop forwarding [21] in Section V. Our numerical results illustrate that the deterministic policy performs close to the complete information optimal closed loop policy for a wide range of parameter values (Section VI).

In a nutshell, the ODE approach is quite common in the modeling of such problems. Its validity in situations without control is established by Kurtz [14], Darling and Norris [16], etc. We aim in this paper at rigorously showing the validity of this limit under control in a few DTN problems.

II. THE SYSTEM MODEL

We consider a set of $K := M + N$ mobile nodes. These include M destinations and N relays. At $t = 0$, a packet is generated and immediately copied to N_0 relays (e.g., via a broadcast from an infrastructure network). Alternatively, these N_0 nodes can be thought of as source nodes.

1) *Mobility model*: We model the point process of the *meeting instants* between pairs of nodes as independent Poisson point processes, each with rate λ . Groenevelt et al. [6] validate this model for a number of common mobility models (random walker, random direction, random waypoint). In particular, they establish its accuracy under the assumptions of small communication range and sufficiently high speed of nodes.

Remarks 2.1: A few studies suggest that traces collected from real-life mobility often demonstrate inter-contact times with power-law distributions. However, Karagiannis et al. [22]

have established that the inter-contact times exhibit exponential tails beyond a certain characteristic time. They also validate this finding across a diverse set of mobility traces. The exponential decay beyond the characteristic time is of relevance as available data traces suggest that the mean inter-contact time is in many cases of the same order as this characteristic time.

2) *Communication model*: Two nodes may communicate only when they come within transmission range of each other, i.e., at *meeting instants*. The transmissions are assumed to be instantaneous. We assume that each transmission of the packet incurs unit energy expenditure at the transmitter.

3) *Relaying model*: We assume that a controlled epidemic relay protocol is employed.

Throughout, we use the terminology relating to the spread of infectious diseases. A node with a copy of the packet is said to be *infected*. A node is said to be *susceptible* until it receives a copy of the packet from another infected node. Thus at $t = 0$, N_0 nodes are infected while $M + N - N_0$ are susceptible.

A. The Forwarding Problem

Our goal is to disseminate the packet to all the M destinations while minimizing the duration until a fraction α ($\alpha < 1$) of the destinations receive the packet.¹

At each meeting epoch with a susceptible relay, an infected node (relay or destination) has to decide whether to copy the packet to the susceptible relay or not. Copying the packet incurs unit cost, but promotes early delivery of the packet to the destinations. We wish to find the trade-off between these costs by minimizing

$$\mathbb{E}\{\mathcal{T}_d + \gamma\mathcal{E}_c\} \quad (1)$$

where \mathcal{T}_d is the time until which at least $M_\alpha := \lceil \alpha M \rceil$ destinations receive the packet, \mathcal{E}_c is the total number of copies made to relays, and γ is the parameter that relates the number of copies to delay cost. Varying γ helps studying the trade-off between the delay and the copying costs.

Remarks 2.2: Alternatively, we may interpret γ as a parameter that accounts for the cost of making a copy, and also relates this cost to the delivery delay. The cost could include the energy cost of transmission and reception, the cost of storage, the price charged by a receiver for carrying a message, etc. Then $\gamma\mathcal{E}_c$ represents the scaled (for comparison with delay) total cost of copying to relays.

Remarks 2.3: Copying the packet to the destinations also incurs a cost. However, this cost is fixed irrespective of the forwarding policy, and thus, is not included in our objective function.

III. OPTIMAL EPIDEMIC FORWARDING

We derive the optimal forwarding policy under the assumption that, at any instant of time, all the nodes have full information about the number of relays carrying the packet and the number of destinations that have received the packet. This assumption will be relaxed in the next section.

¹Subsequently, we analyze a scaled version of the network in order to obtain a distributed policy. See Footnote 9 for why we restrict to $\alpha < 1$.

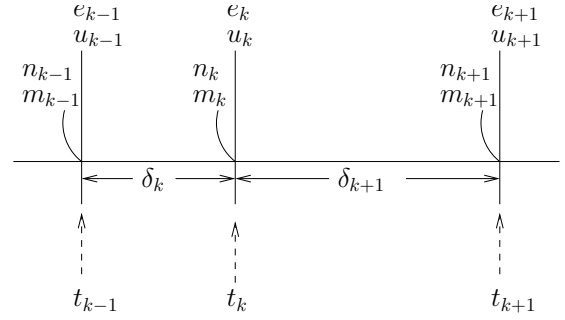


Fig. 1. Evolution of the controlled Markov chain $\{s_k\}$. Note that (m_k, n_k) is embedded at t_k- , i.e., just before the meeting epoch.

A. The MDP Formulation

Let $t_0 := 0$ and $t_k, k = 1, 2, \dots$ denote the meeting epochs of the infected nodes (relays or destinations) with the susceptible nodes. Define $\delta_k := t_k - t_{k-1}$ for $k \geq 1$. Let $m(t)$ and $n(t)$ be the numbers of infected destinations and relays, respectively, at time t . Thus $m(0) = 0$ and $n(0) = N_0$, and the forwarding process stops at time t if $m(t) = M$. We use m_k and n_k to mean $m(t_k-)$ and $n(t_k-)$ which are the numbers of infected destinations and relays, respectively, just before the meeting epoch t_k . Let e_k be the type of the susceptible node that an infected node meets at t_k ; $e_k \in \mathcal{E} := \{d, r\}$ where d and r stand for destination and relay, respectively. The state of the system at a meeting epoch t_k is given by the tuple

$$s_k := (m_k, n_k, e_k).$$

Since the forwarding process stops at time t if $m(t) = M$, the state space is $[M - 1] \times [N_0 : N] \times \mathcal{E}$.²

Let u_k be the action of the infected node at meeting epoch $t_k, k = 1, 2, \dots$. The control space is $\mathcal{U} \in \{0, 1\}$, where 1 is for *copy* and 0 is for *do not copy*. The embedding convention described above is shown in Figure 1.

We treat the tuple (δ_{k+1}, e_{k+1}) as the random disturbance at epoch t_k . Note that for $k = 1, 2, \dots$, the time between successive decision epochs, δ_k , is independent and exponentially distributed with parameter $(m_k + n_k)(M + N - m_k - n_k)\lambda$. Furthermore, with “w.p.” standing for “with probability”,

$$e_k = \begin{cases} d & \text{w.p. } p_{m_k, n_k}(d) := \frac{M - m_k}{M + N - m_k - n_k}, \\ r & \text{w.p. } p_{m_k, n_k}(r) := \frac{N - n_k}{M + N - m_k - n_k}. \end{cases}$$

1) *Transition structure*: From the description of the system model, the state at time $k + 1$ is $s_{k+1} = (m_k + u_k, n_k, e_{k+1})$ if $e_k = d$, and $s_{k+1} = (m_k, n_k + u_k, e_{k+1})$ if $e_k = r$. Recall that e_{k+1} is a component in the random disturbance. Thus the next state is a function of the current state, the current action and the current disturbance as required for an MDP.

2) *Cost Structure*: For a state-action pair (s_k, u_k) the expected single stage cost is given by

$$g(s_k, u_k) = \gamma u_k 1_{\{e_k=r\}} + \mathbb{E}\{\delta_{k+1} 1_{\{m_{k+1} < M_\alpha\}}\},$$

where the expectation is taken with respect to the random disturbance (δ_{k+1}, e_{k+1}) . It can be observed that

²We use notation $[a] = \{0, 1, \dots, a\}$ and $[a : b] = \{a, a + 1, \dots, b\}$ for $b \geq a + 1$ and $a, b \in \mathbb{Z}_+$.

$$g(s_k, u_k) = \begin{cases} \gamma u_k 1_{\{e_k=r\}} & \text{if } m_k \geq M_\alpha, \\ 0 & \text{if } m_k = M_\alpha - 1, e_k = d \text{ and } u_k = 1, \\ \gamma u_k 1_{\{e_k=r\}} + C_d(s_k, u_k) & \text{otherwise,} \end{cases}$$

where

$$C_d(s_k, u_k) = \frac{1}{(m_k + n_k + u_k)(M + N - m_k - n_k - u_k)\lambda}$$

is the mean time until the next decision epoch. The quantity γ is expended whenever $u_k = 1$, i.e., the action is to copy.

3) *Policies*: A policy π is a sequence of mappings $\{u_k^\pi, k = 0, 1, 2, \dots\}$, where $u_k^\pi : [M-1] \times [N_0 : N] \times \mathcal{E} \rightarrow \mathcal{U}$. The cost of a policy π for an initial state $s = (m, n, e)$ is

$$J_\pi(s) = \sum_{k=0}^{\infty} \mathbb{E} \left\{ g(s_k, u_k^\pi(s_k)) \mid s_0 = s \right\}.$$

Let Π be the set of all policies. Then the optimal cost function is defined as

$$J(s) = \min_{\pi \in \Pi} J_\pi(s).$$

A policy π is called stationary if u_k^π are identical, say u , for all k . For brevity we refer to such a policy as the stationary policy u . A stationary policy $u^* \equiv \{u^*, u^*, \dots\}$ is optimal if $J_{u^*}(s) = J(s)$ for all states s .

4) *Total Cost*: We now translate the optimal cost-to-go from the first decision instant t_1 into the optimal total cost. Recall that at t_1 , the state s_1 is $(0, N_0, r)$ or $(0, N_0, d)$ depending on whether the susceptible node that is met is a relay or a destination. The objective function (1) can then be restated as

$$\mathbb{E}_\pi \{ \mathcal{T}_d + \gamma \mathcal{E}_c \} = \frac{1}{\lambda N_0 (M + N - N_0)} + \left(\frac{N - N_0}{M + N - N_0} J_\pi(0, N_0, r) + \frac{M}{M + N - N_0} J_\pi(0, N_0, d) \right), \quad (2)$$

where the subscript π shows dependence on the underlying policy. In the right hand side, the first term $\frac{1}{\lambda N_0 (M + N - N_0)}$ is the average delay until the first decision instant which has to be borne under any policy.

B. Optimal Policy

Since the cost function $g(\cdot)$ is nonnegative, Proposition 1.1 in [19, Chapter 3] implies that the optimal cost function will satisfy the following Bellman equation. For $s = (m, n, e)$,

$$J(s) = \min_{u \in \{0,1\}} A(s, u)$$

$$\text{where } A(s, u) = g(s, u) + \mathbb{E}(J(s') \mid s, u).$$

Here s' denotes the next state which depends on s, u and the random disturbance in accordance with the transition structure described above. The expectation is taken with respect to the random disturbance. Since the action space is finite, there exists a stationary policy u^* such that, for all s , $u^*(s)$ attains minimum in the above Bellman equation (see [19, Chapter 3]). Now we characterize this stationary optimal policy.

First, observe that it is optimal to copy whenever a susceptible destination is encountered; it does not incur any cost

and increases the number of infected nodes, which in turn also expedites the packet delivery process. Moreover, every destination has to be infected at some time. Thus $u^*(m, n, d) = 1$ for all $(m, n) \in [M-1] \times [N_0 : N]$. Next, once M_α destinations have been infected, no further delay cost is incurred, and so further copying to relays does not help. Thus $u^*(m, n, r) = 0$ for all $(m, n) \in [M_\alpha : M-1] \times [N_0 : N]$.

Focus now on a reduced state space $[M_\alpha - 1] \times [N_0 : N] \times \{r\}$. Consider the following *one step look ahead policy* [19, Section 3.4]. At a meeting with a susceptible relay, when the state is (m, n, r) , compare the following two action sequences.

- 1) 0s: *stop*, i.e., do not copy to this relay or to any susceptible relays met in the future,
- 2) 1s: copy to this relay and then *stop*.

The costs to go corresponding to the action sequences 0s and 1s are, respectively,

$$J_{0s}(m, n, r) = \sum_{j=m}^{M_\alpha-1} \frac{1}{\lambda(n+j)(M-j)}$$

$$\text{and } J_{1s}(m, n, r) = \gamma + \sum_{j=m}^{M_\alpha-1} \frac{1}{\lambda(n+j+1)(M-j)}.$$

The *stopping set* \mathcal{S}_S is defined to be

$$\mathcal{S}_S := \{(m, n, r) : \Phi(m, n) \leq 0\} \quad (3)$$

where

$$\begin{aligned} \Phi(m, n) &:= J_{0s}(m, n, r) - J_{1s}(m, n, r) \\ &= \sum_{j=m}^{M_\alpha-1} \frac{1}{\lambda(n+j)(n+j+1)(M-j)} - \gamma \end{aligned} \quad (4)$$

for all $(m, n) \in [M_\alpha - 1] \times [N_0 : N]$. The one step look ahead policy is to copy to relay when $(m, n, r) \notin \mathcal{S}_S$, and to stop copying otherwise.³

One step look ahead policies are shown to be optimal for stopping problems under certain conditions (see [23, Section 4.4] and [19, Section 3.4]). But our problem is not a stopping problem as an action 0 is not equivalent to *stop*; even if the susceptible relay met now is not copied, the resulting state is not a *terminal state*, and one met in the future may be copied. However, we exploit the cost structure to prove that when an infected node meets a susceptible relay, it can restrict attention to two actions: 1 (i.e., copy now) and *stop* (i.e., do not copy now and never copy again). Subsequently, we also show that the above one step look ahead policy is optimal.

Theorem 3.1: The optimal policy $u^* : [M-1] \times [N_0 : N] \times \mathcal{E} \rightarrow \mathcal{U}$ satisfies

$$u^*(m, n, e) = \begin{cases} 1, & \text{if } e = d, \\ 1, & \text{if } e = r \text{ and } \Phi(m, n) > 0, \\ \text{stop} & \text{if } e = r \text{ and } \Phi(m, n) \leq 0. \end{cases}$$

Proof: See Appendix A. ■

³Convention: A sum over an empty set is 0. Thus $\Phi(m, n) = -\gamma$ if $m \geq M_\alpha$. Consequently, for the states $[M_\alpha : M-1] \times [N_0 : N] \times \{r\}$, one step-look ahead policy prescribes *stop*. This is consistent with our earlier discussion.




Fig. 2. An illustration of the optimal policy. The symbols 'x' mark the states in which the optimal action (at meeting with a relay) is to copy

Remarks 3.1: We can define $\Phi(m, n)$ also for the case $\alpha = 1$, i.e., when we attempt to minimize the delay until all the destinations receive the packet. Theorem 3.1 continues to hold.

We illustrate the optimal policy using an example. Let $M = 15, N = 50, N_0 = 10, \alpha = 0.8, \lambda = 0.001$ and $\gamma = 1$. The “x” in Figure 2 are the states where the optimal action (at meeting with a relay) is to copy. For example, if only 5 destinations have the packet, then relays are copied to if and only if there are 24 or less infected relays. If 7 destinations already have the packet and there are 19 infected relays, then no further copying to relays is done.

IV. ASYMPTOTICALLY OPTIMAL EPIDEMIC FORWARDING

In states $[M_\alpha - 1] \times [N_0 : N] \times \{r\}$, the optimal action, which is governed by the function $\Phi(m, n)$, requires perfect knowledge of the network state (m, n) . This may not be available to the decision maker due to intermittent connectivity. In this section, we derive an asymptotically optimal policy that does not require knowledge of network's state but depends only on the time elapsed since the generation of the packet. Such a policy is implementable if the packet is time-stamped when generated and the nodes' clocks are synchronized.

A. Asymptotic Deterministic Dynamics

Our analysis closely follows Darling [15]. It is straightforward to show that the equations that follow are the conditional expected drift rates of the optimally controlled CTMC. For $(m(t), n(t)) \in [M - 1] \times [N_0 : N]$, using the optimal policy in Theorem 3.1, we get

$$\frac{d\mathbb{E}(m(t)|(m(t), n(t)))}{dt} = \lambda(m(t) + n(t))(M - m(t)), \quad (5a)$$

$$\frac{d\mathbb{E}(n(t)|(m(t), n(t)))}{dt} = \lambda(m(t) + n(t))(N - n(t)) \mathbf{1}_{\{\Phi(m(t), n(t)) > 0\}}. \quad (5b)$$

Recalling that $K = M + N$, the total number of nodes, we study large K asymptotics. More precisely, we consider

a sequence of problems with increasing M, N, N_0 (and thus also $K := M + N$) such that the ratios $\frac{M}{K}, \frac{N}{K}$ and $\frac{N_0}{K}$ remain constant. The problems are indexed by K . The parameters of the K th problem are denoted using the superscript K . Normalized versions of these parameters, and normalized versions of the system state are denoted as follows:

$$\left. \begin{aligned} X &= \frac{M^K}{K}, Y = \frac{N^K}{K}, \\ X_\alpha &= \frac{\alpha M^K}{K}, Y_0 = \frac{N_0^K}{K}, \\ \lambda^K &= \frac{\Lambda}{K}, \gamma^K = \frac{\Gamma}{K}, \\ x^K(t) &= \frac{m^K(t)}{K} \text{ and } y^K(t) = \frac{n^K(t)}{K}. \end{aligned} \right\} \quad (6)$$

Remarks 4.1: The pairwise meeting rate and the copying cost must both scale down as K increases. Otherwise, the delivery delay will be negligible and the total transmission cost will be enormous for any policy, and no meaningful analysis is possible.

For each K , we define scaled two-dimensional integer lattice

$$\Delta^K = \left\{ \left(\frac{i}{K}, \frac{j}{K} \right) : (i, j) \in [M^K - 1] \times [N_0^K : N^K] \right\}.$$

Clearly, $(x^K(t), y^K(t)) \in \Delta^K$. Now, using the notation in (6), the drift rates in (5a)-(5b) can be rewritten as follows.

$$\begin{aligned} \frac{d\mathbb{E}(x^K(t)|(x^K(t), y^K(t)))}{dt} &= f_1^K(x^K(t), y^K(t)) \\ &:= \Lambda(x^K(t) + y^K(t))(X - x^K(t)), \end{aligned} \quad (7a)$$

$$\begin{aligned} \frac{d\mathbb{E}(y^K(t)|(x^K(t), y^K(t)))}{dt} &= f_2^K(x^K(t), y^K(t)) \\ &:= \Lambda(x^K(t) + y^K(t))(Y - y^K(t)) \mathbf{1}_{\{\phi^K(x^K(t), y^K(t)) > 0\}}, \end{aligned} \quad (7b)$$

where, for $(x, y) \in \Delta^K$,

$$\phi^K(x, y) := \sum_{j=Kx}^{\lceil KX_\alpha \rceil - 1} \frac{1}{K\Lambda(y + \frac{j}{K})(y + \frac{j+1}{K})(X - \frac{j}{K})} - \Gamma. \quad (8)$$

We also define $(x(t), y(t)) \in [0, X] \times [Y_0, Y]$ as functions satisfying the following ODEs: $x(0) = 0, y(0) = Y_0$, and for $t \geq 0$,

$$\frac{dx(t)}{dt} = f_1(x(t), y(t)) := \Lambda(x(t) + y(t))(X - x(t)), \quad (9a)$$

$$\frac{dy(t)}{dt} = f_2(x(t), y(t)) := \Lambda(x(t) + y(t))(Y - y(t)) \mathbf{1}_{\{\phi(x(t), y(t)) > 0\}} \quad (9b)$$

where⁴

$$\phi(x, y) = \int_{z=x}^{X_\alpha} \frac{dz}{\Lambda(y+z)^2(X-z)} - \Gamma. \quad (10)$$

⁴Convention: An integral assumes the value 0 if its lower limit exceeds the upper limit. So, $\phi(x, y) = -\Gamma$ if $x \geq X_\alpha$.

Finally, we redefine the delivery delay \mathcal{T}_d (see (1)) to be

$$\tau^K = \inf\{t \geq 0 : x^K(t) \geq X_\alpha\}, \quad (11)$$

$$\text{and } \tau = \inf\{t \geq 0 : x(t) \geq X_\alpha\}. \quad (12)$$

Note that τ^K is a stopping time for the random process $(x^K(t), y^K(t))$, whereas τ is a deterministic time instant. Since $f_1^K(x, y)$ is bounded away from zero, $\tau^K < \infty$ with probability 1. Similarly, on account of $f_1(x, y)$ being bounded away from zero, $\tau < \infty$.

Kurtz [14] and Darling [15] studied convergence of CTMCs to the solutions of ODEs. The following are the hypotheses for the version of the limit theorem that appears in Darling [15].

- (i) $\lim_{K \rightarrow \infty} \mathbb{P}(\|(x^K(0), y^K(0)) - (x(0), y(0))\| > \epsilon) = 0$;
- (ii) In the scaled process $(x^K(t), y^K(t))$, the jump rates are $O(K)$ and drifts are $O(K^{-1})$;
- (iii) $(f_1^K(x, y), f_2^K(x, y))$ converges to $(f_1(x, y), f_2(x, y))$ uniformly in (x, y) ;
- (iv) $(f_1(x, y), f_2(x, y))$ is Lipschitz continuous.

Observe that, in our case, only the first two hypotheses are satisfied. In particular, $f_2^K(x, y)$ does not converge uniformly to $f_2(x, y)$, and $f_2(x, y)$ is not Lipschitz over $[0, X_\alpha] \times [Y_0, Y]$. Hence, the convergence results do not directly apply in our context. Thankfully, there is some regularity we can exploit which we now summarize as easily verifiable facts.

- (a) $\phi^K(x, y)$ converges uniformly to $\phi(x, y)$;
- (b) the drift rates $f_1(x, y)$ and $f_2(x, y)$ are bounded from below and above;
- (c) $f_1(x, y)$ is Lipschitz and $f_2(x, y)$ is locally Lipschitz; and
- (d) for all small enough $\nu \in \mathbb{R}$, and all (x, y) on the graph of “ $\phi(x, y) = \nu$ ”, the direction in which the ODE progresses, $(f_1(x, y), f_2(x, y))$, is not tangent to the graph.

We then prove the following result which is identical to [15, Theorem 2.8].

Theorem 4.1: Assume that $\alpha < 1$ and $Y_0 > 0$. Then, for every $\epsilon, \delta > 0$, for the optimally controlled epidemic forwarding process,

$$\lim_{K \rightarrow \infty} \mathbb{P}\left(\sup_{0 \leq t \leq \tau} \|(x^K(t), y^K(t)) - (x(t), y(t))\| > \epsilon\right) = 0,$$

$$\lim_{K \rightarrow \infty} \mathbb{P}(|\tau^K - \tau| > \delta) = 0.$$

Proof: See Appendix B. ■

We illustrate Theorem 4.1 using an example. Let $X = 0.2, Y = 0.8, \alpha = 0.8, Y_0 = 0.2, \Lambda = 0.05$ and $\Gamma = 50$. In Figure 3, we plot $(x(t), y(t))$ and sample trajectories of $(x^K(t), y^K(t))$ for $K = 100, 200$ and 500 . We indicate the states at which the optimal policy stops copying to relays, i.e., $\Phi^K(x^K(t), y^K(t))$ goes below 0 (see Theorem 3.1) and the states at which the fraction of infected destinations crosses X_α . We also show the corresponding states in the fluid model. The plots show that for large K , the fluid model captures the random dynamics of the network very well.

B. Asymptotically Optimal Policy

Observe that $\phi(x, y)$ is decreasing in x and y , both of which are nondecreasing with t . Consequently $\phi(x(t), y(t))$




Fig. 3. Simulation results: The top and bottom subfigures respectively show the fractions of infected destinations and relays as a function of time. $(x^K(t), y^K(t))$ are obtained from a simulation of the controlled CTMC, and $(x(t), y(t))$ from the ODEs. The marker ‘X’ indicates the states at which copying to relays is stopped whereas ‘O’ indicates the states at which a fraction α of destinations have the packet.

decreases with t . We define

$$\tau^* := \inf\{t \geq 0 : \phi(x(t), y(t)) \leq 0\}. \quad (13)$$

The limiting deterministic dynamics suggests the following policy u^∞ for the original forwarding problem.⁵

$$u^\infty(m, n, e) = \begin{cases} 1 & \text{if } e = d, \\ 1 & \text{if } e = r \text{ and } t \leq \tau^*, \\ 0 & \text{if } e = r \text{ and } t > \tau^*. \end{cases}$$

We show that the policy u^∞ is asymptotically optimal in the sense that its expected cost approaches the expected cost of the optimal policy u^* as the network grows. Let us restate (2) as

$$\mathbb{E}_\pi^K \{\mathcal{T}_d + \gamma \mathcal{E}_c\} = \frac{1}{K \Lambda Y_0 (1 - Y_0)} + \left(\frac{Y - Y_0}{1 - Y_0} J_\pi^K(0, Y_0, r) + \frac{X}{1 - Y_0} J_\pi^K(0, Y_0, d) \right).$$

We have used superscript K to show the dependence of cost on the network size. We then establish the following asymptotic optimality result.

⁵Observe that the policy u^∞ does not require knowledge of m and n . The infected node readily knows the type of the susceptible node (d or r) at the decision epoch.

Theorem 4.2: Assume that $\alpha < 1$ and $Y_0 > 0$. Then

$$\lim_{K \rightarrow \infty} \mathbb{E}_{u^*}^K \{\mathcal{T}_d + \gamma \mathcal{E}_c\} = \lim_{K \rightarrow \infty} \mathbb{E}_{u^\infty}^K \{\mathcal{T}_d + \gamma \mathcal{E}_c\} = \tau + \Gamma y(\tau^*).$$

Proof: See Appendix C. ■

Remarks 4.2: Observe that we do not compare the limiting value of the optimal cost with the optimal cost on the (limiting) deterministic system. In general, these two may differ.⁶ However, the deterministic policy u^∞ can be applied on the finite K -node system. The above theorem asserts that given any $\epsilon > 0$, cost of the policy u^∞ is within ϵ of the optimal cost on the K -node system for all sufficiently large K .

Distributed Implementation: The asymptotically optimal policy can be implemented in a distributed fashion. We assume that the system parameters $M, N, \alpha, N_0, \lambda$ and γ are known at the source, and also that all the nodes are time synchronized.⁷ Suppose that the packet is generated at the source at time t_0 (we assumed $t_0 = 0$ for the purpose of analysis). Given the system parameters, the source first extracts $X, Y, X_\alpha, Y_0, \Lambda$ and Γ as in (6). Then, it calculates τ^* (see (13)), and stores $t_0 + \tau^*$ as a header in the packet.

The packet is immediately copied to N_0 relays, perhaps by means of a broadcast from an infrastructure “base station”. When an infected node meets a susceptible relay, it compares $t_0 + \tau^*$ with the current time. The susceptible relay is not copied to if the current time exceeds $t_0 + \tau^*$; the nodes do not need to know the transient numbers of infected relays and infected destinations. However, all the infected nodes continue to carry the packet, and to copy to susceptible destinations as and when they meet.

Remarks 4.3: Consider a scenario, where the interest is in copying packet to only a fraction α of the destinations. Observe from Theorem 4.1 that for every $\epsilon > 0$,

$$\lim_{K \rightarrow \infty} \mathbb{P} \left(\left| \frac{m^K(\tau)}{M} - \alpha \right| > \epsilon \right) = 0.$$

Thus, in large networks, copying to destinations can also be stopped at time τ (see (12)) while ensuring that with large probability the fraction of infected destinations is close to α . Consequently, all the relays can delete the packet and free their memory at τ . This helps when packets are large and relay (cache) memory is limited.

V. OPTIMAL TWO-HOP FORWARDING

Instead of epidemic relaying one can consider two-hop relaying [21]. Here, the N_0 source nodes can copy the packet to any of the $N - N_0$ relays or M destinations. The infected destinations can also copy the packet to any of the susceptible relays or destinations. However, the relays are allowed to transmit the packet only to the destinations. Here also a similar optimization problem as in Section II-A arises.

⁶In our case these two indeed match. See [24, Appendix D] for a proof.

⁷In practice, due to variations in the clock frequency, the clocks at different nodes will drift from each other. But the time differences are negligible compared to the delays caused by intermittent connectivity in the network. Moreover, when an infected node meets a susceptible node, clock synchronization can be performed before the packet is copied. Distributed algorithms for time synchronization are also available (e.g., see [25]).

Now, the decision epochs $t_k, k = 1, 2, \dots$ are the meeting epochs of the infected nodes (sources, relays or destinations) with the susceptible destinations and the meeting epochs of the sources or infected destinations with the susceptible relays. We can formulate an MDP with state

$$s_k := (m_k, n_k, e_k).$$

at instant t_k where m_k, n_k and e_k are as defined in Section III-A. The state space is $[M_\alpha - 1] \times [N_0 : N] \times \mathcal{E}$. The control space is $\mathcal{U} \in \{0, 1\}$, where 1 is for *copy* and 0 is for *do not copy*. We also get a transition structure identical to that in Section III-A.

For a state action pair (s_k, u_k) the expected single stage cost is given by

$$g(s_k, u_k) = \gamma u_k 1_{\{e_k=r\}} + \mathbb{E} \left\{ \delta_{k+1} 1_{\{m_{k+1} < M_\alpha\}} \right\} \\ = \begin{cases} \gamma u_k 1_{\{e_k=r\}} & \text{if } m_k \geq M_\alpha, \\ 0 & \text{if } m_k = M_\alpha - 1, e_k = d \text{ and } u_k = 1, \\ \gamma u_k 1_{\{e_k=r\}} + C_d(s_k, u_k) & \text{otherwise,} \end{cases}$$

where

$$C_d(s_k, u_k) = ((m_k + n_k + u_k)(M - m_k - u_k 1_{\{s_k=d\}})\lambda \\ + (m_k + u_k 1_{\{s_k=d\}} + N_0)(N - n_k - u_k 1_{\{s_k=r\}})\lambda)^{-1}$$

is the mean time until the next decision epoch. As before, the quantity γu_k accounts for the transmission energy.

Let $u^* : [M_\alpha - 1] \times [N_0 : N] \times \mathcal{E} \rightarrow \mathcal{U}$ be a stationary optimal policy. As in Section III-B, the optimal policy satisfies $u^*(m, n, d) = 1$ for all $(m, n) \in [M - 1] \times [N_0 : N]$, and $u^*(m, n, r) = 0$ for all $(m, n) \in [M_\alpha : M - 1] \times [N_0 : N]$. Thus, we focus on a reduced state space $[M_\alpha - 1] \times [N_0 : N] \times \{r\}$. As before, we look for the one step look ahead policy which turns out to be the same as that for epidemic relaying. Finally, Theorem 3.1 holds for two-hop relaying as well (see the proof in Appendix A).

Next, we turn to the asymptotically optimal control for two-hop relaying. The following are the conditional expected drift rates. For $(m(t), n(t)) \in [M_\alpha - 1] \times [N_0 : N]$,

$$\frac{d\mathbb{E}(m(t)|(m(t), n(t)))}{dt} = \lambda(m(t) + n(t))(M - m(t)), \\ \frac{d\mathbb{E}(n(t)|(m(t), n(t)))}{dt} = \lambda(m(t) + N_0)(N - n(t)) \\ 1_{\{\Phi(m(t), n(t)) > 0\}}.$$

We employ the same scaling and notations as in (6). The drift rates in terms of $(x^K(t), y^K(t)) \in [0, X_\alpha] \times [Y_0, Y]$ are

$$\frac{d\mathbb{E}(x^K(t)|(x^K(t), y^K(t)))}{dt} = f_1^K(x^K(t), y^K(t)) \\ := \Lambda(x^K(t) + y^K(t))(X - x^K(t)), \\ \frac{d\mathbb{E}(y^K(t)|(x^K(t), y^K(t)))}{dt} = f_2^K(x^K(t), y^K(t)) \\ := \Lambda(x^K(t) + Y_0)(Y - y^K(t)) 1_{\{\phi^K(x^K(t), y^K(t)) > 0\}},$$

Now, $x(t), y(t)$ are defined as functions satisfying $x(0) =$

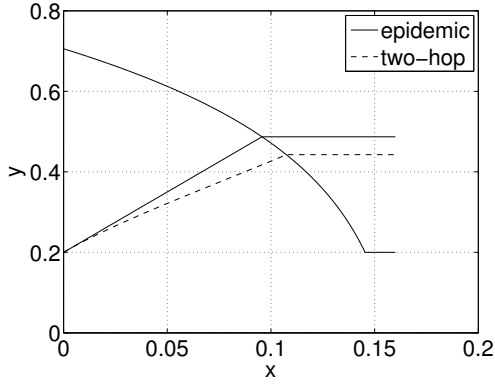


Fig. 4. An illustration of the epidemic and two hop trajectories. The plots also show the graph of ‘ $\phi(x, y) = 0$ ’.

$0, y(0) = Y_0$ and for $t \geq 0$,

$$\begin{aligned} \frac{dx(t)}{dt} &= f_1(x(t), y(t)) := \Lambda(x(t) + y(t))(X - x(t)), \\ \frac{dy(t)}{dt} &= f_2(x(t), y(t)) := \Lambda(x(t) + Y_0)(Y - y(t)) \\ &\quad \mathbb{1}_{\{\phi(x(t), y(t)) > 0\}} \end{aligned}$$

The analysis in Section IV applies to two-hop relaying as well. In particular, Theorems 4.1 and 4.2 hold. However, for the identical system parameters (M, N, α, λ and γ) and initial state (N_0), the value of the time-threshold τ^* will be larger due to the slower rates of infection of relays and destinations.

We illustrate the comparison between epidemic and two-hop relaying via an example. Let $X = 0.2, Y = 0.8, \alpha = 0.8, Y_0 = 0.2, \Lambda = 0.05$ and $\Gamma = 50$. In Figure 4, we plot the graph of “ $\phi(x, y) = 0$ ”, and also the ‘ y versus x ’ trajectories corresponding to epidemic and two-hop relaying. In Figure 5, we plot the trajectories of $x(t)$ and $y(t)$ corresponding to epidemic and two-hop relaying. As anticipated, the value of the time-threshold τ^* is larger for two-hop relaying than epidemic relaying. Moreover, the number of transmissions is less while the delivery delay is more under the controlled two-hop relaying.

VI. SIMULATION AND NUMERICAL RESULTS

We start with simulations that validate the independent Poisson point process model for the meeting instants in the presence of control. We then show a few numerical results to demonstrate the good performance of the deterministic control, and also to compare the performance of optimal epidemic and two-hop relaying.

For simulation we use the mobility traces generated using OMNeT++ according to the *random waypoint* mobility model [6]. In our setup $K = 200$ nodes move at a speed 10 m/s in a square of size 2 Km \times 2 Km, and have a communication radius 28.7115 m. These values are chosen to yield a pair-wise meeting rate $\lambda = 0.00025/s$. We take $M = 20, N = 180, \alpha = 0.8$ and $N_0 = 20$, and assume quarter a unit cost per copy ($\gamma = 0.25$). We perform 1000 runs of the simulation, and average $m(t)$ and $n(t)$; averaging leads to fractional values of these otherwise integer variables. We choose destinations and




Fig. 5. The top and bottom subfigures respectively show the fractions of infected destinations and relays as a function of time. The marker ‘X’ indicates the states at which copying to relays is stopped (i.e., the states at τ^*), and ‘O’ indicates the states at which α fraction of destinations have been copied (i.e., the states at τ).

initially infected relays uniformly at random from among all the nodes for each iteration. We consider both epidemic and two-hop relaying. As Figure 6 depicts, in both the cases, the Poisson point process model for the meeting instants quite accurately predicts the dynamics of the fraction of infected destinations and relays under optimal control.

Now, we use the Poisson point process model for the meeting instants to illustrate the performance of optimal and deterministic open loop controls. We set $M = 20, N = 80, \alpha = 0.8, N_0 = 10$ and $\lambda = 0.0005$, and vary γ from 0.01 to 10. In Figure 7, we plot the total number of copies to relays, delivery delays and total costs. The subfigures on the left are for epidemic relaying while those on the right are for two-hop relaying. Each subfigure contains four plots: the optimal policy, the deterministic open loop policy, the spray and wait policy [9], and uncontrolled forwarding policy (copying to all the relays).⁸ The authors in [9] propose *spray and wait* and *spray and focus* routings for single destination networks, and select the number of copies to meet certain expected delay target. We have adapted these for a multi-destination network, and we choose the number of copies to minimize the weighted sum of delivery delay and copying cost as in (1). In the *spray* phase, we use *binary spraying* in the left hand

⁸In our symmetric Poisson meeting model, *spray and wait* and *spray and focus* routings are identical; no relays are copied in the *focus* phase.

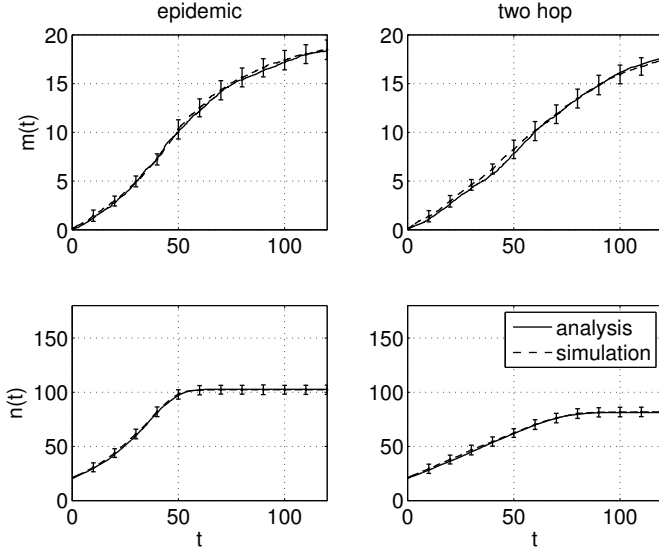


Fig. 6. The top and bottom subfigures, respectively, show the numbers of infected destinations and relays as a function of time. The simulation plots are based on traces generated according to random waypoint mobility, while the analysis plots are based on the Poisson point process model for the meeting instants. The left plots are for optimally controlled epidemic relaying and the right plots are for optimally controlled two-hop relaying. We also show 95% confidence intervals for our simulation.

side plots and *source spraying* in the right hand side plots. Clearly, the deterministic policy performs close to the optimal policy for all the considered parameter sets. Both these policies outperform *spray and wait* routing as well as uncontrolled forwarding. Performance improvements with respect to *spray and wait* routing are substantial for small values of γ . On the other hand, performance improvements with respect to the uncontrolled forwarding are meager for small values of γ ; for small copying costs the controlled protocols also make a large number of copies (because the difference in the number of copies affects the total cost only marginally), and hence the controlled protocols incur approximately the similar delivery delays as the uncontrolled one. However, for higher values of γ , optimal control brings in considerable performance benefits.

Often static (probabilistic) controls have been considered in the literature (e.g., see [10]). In our context, probabilistic controls end up copying to all the relays, but incur higher delivery delay than uncontrolled forwarding. In particular, such controls incur higher total cost than uncontrolled forwarding, and are clearly suboptimal.

Our results also characterize a few typical features of epidemic and two-hop relayings. As expected, for each set of parameters, the optimal number of copies is less and the optimal delivery delay is more under two-hop relaying when compared with epidemic relaying. We also observe that the optimal epidemic relaying always performs better than the optimal two-hop relaying. This also is expected because in epidemic relaying we optimally control forwarding without the constraint that relays cannot copy to other susceptible relays. However, for large values of γ , controlled epidemic as well as controlled two-hop relayings provision few copies to relays. Thus, in this regime, both these schemes incur almost same

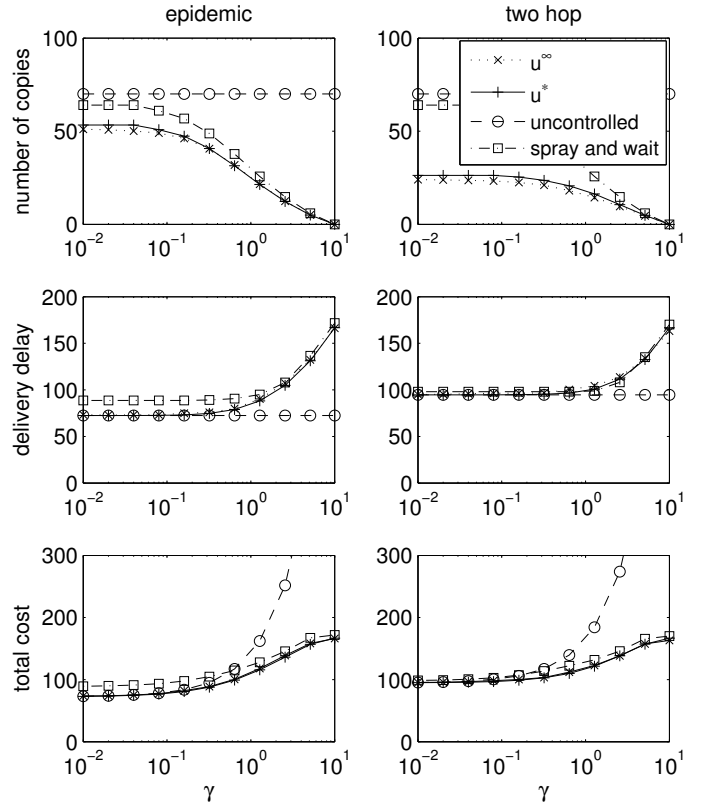


Fig. 7. The top, middle and bottom subfigures, respectively, show expected total number of copies to relays, expected delivery delays and expected total costs corresponding to the optimal, the deterministic and *spray and wait* policies, and also corresponding to uncontrolled forwarding. The left plots are for epidemic relaying and the right plots are for two-hop relaying.

delivery delays, and so incur almost same total costs.

Now, we investigate the effect of α on the number of copies made and the delivery delay. Towards this, we set $M = 40$, $N = 160$, $N_0 = 20$, $\lambda = 0.00001$ and $\gamma = 2$, and vary α from 0.02 to 0.97 such that M_α assumes all the integral values from 1 to 39. We focus only on epidemic relaying. Expectedly, both the number of copies and the delivery delay increase with α (see Figure 8). In particular, as α approaches one, even the optimal policy copies to almost all the susceptible relays (in this example, 140), and all the policies (including uncontrolled forwarding) incur huge delivery delays.⁹

Finally, we study the effect of varying the network size K and the pair-wise meeting rate λ . Let $X = 0.2$, $Y = 0.8$, $\alpha = 0.8$, $Y_0 = 0.1$ and $\gamma = 0.1$. We vary λ from 0.0001 to 0.1 and use $K = 50, 100$ and 200 . In Figure 9, we plot the number of copies to relays and the delivery delays corresponding to both the optimal and the asymptotically optimal deterministic policies in the case of epidemic relaying. We observe that, for a fixed K , both the mean delivery delay and the mean number of copies decrease as λ increases. We also observe that, for

⁹ As α approaches one, $\phi(x, y)$ approaches infinity (see (10)), and consequently τ^* also approaches infinity (see (13)). Thus, the deterministic policy prescribes copying to all the relays in an attempt to mitigate the enormous delivery delay. This illustrates why we restrict to $\alpha < 1$ to get a meaningful distributed policy.

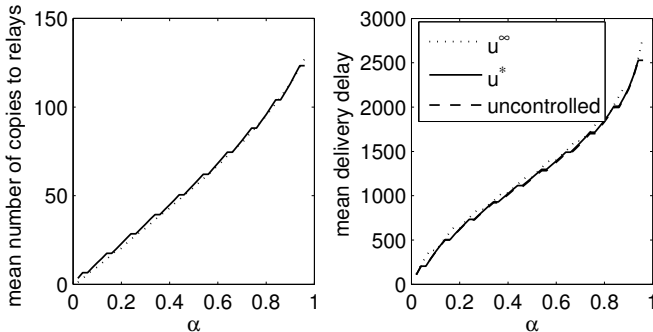


Fig. 8. The left and right subfigures, respectively, show the expected total number of copies to relays and expected delivery delays corresponding to both the optimal and the deterministic policies under controlled epidemic relaying. The mean delivery delay under uncontrolled epidemic relaying is also shown.

a fixed λ , the mean delivery delay decreases as the network size grows. Finally, for smaller values of λ , the mean number of copies to relays increases with the network size, and for larger values of λ , the opposite happens.

VII. CONCLUSION

We studied the epidemic forwarding in DTNs, formulated the problem as a controlled continuous time Markov chain, and obtained the optimal policy (Theorem 3.1). We then developed an ordinary differential equation approximation for the optimally controlled Markov chain, under a natural scaling, as the population of nodes increases to ∞ (Theorem 4.1). This o.d.e. approximation yielded a forwarding policy that does not require global state information (and, hence, is implementable), and is asymptotically optimal (Theorem 4.2).

The optimal forwarding problem can also be addressed following the result of Gast et al. [18]. They study a general discrete time Markov decision process (MDP) [19]. However, they do not solve the finite problem citing the difficulties associated with obtaining the asymptotics of the optimally controlled process (see [18, Section 3.3]). Instead, they consider the fluid limit of the MDP, and analyze optimal control over the deterministic limiting problem. They then show that the optimal reward of the MDP converges to the optimal reward of its mean field approximation, given by the solution to a Hamilton-Jacobi-Bellman (HJB) equation [23, Section 3.2]. On the other hand, our approach is more direct. We have a continuous time controlled Markov chain at our disposal. We explicitly characterize the optimal policy for the finite (complete information) problem, and prove convergence of the optimally controlled Markov chain to a fluid limit. An asymptotically optimal deterministic control is then suggested by the limiting deterministic dynamics, and does not require solving HJB equations. Our notion of asymptotic optimality is also stronger in the sense that we apply both the optimal policy and the deterministic policy to the finite problem, and show that the corresponding costs converge.

There are several directions in which this work can be extended. In the same DTN framework, there could be a deadline on the delivery time of the packet (or message); the goal of the optimal control could be to maximize the fraction

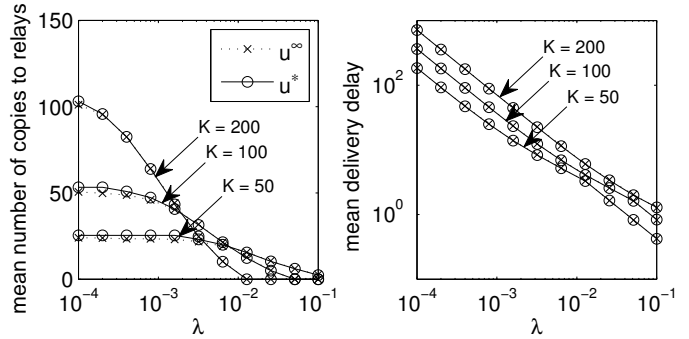


Fig. 9. The left and right subfigures, respectively, show the expected total number of copies to relays and expected delivery delays corresponding to both the optimal and the deterministic policies under controlled epidemic relaying.

of destinations that receive the packet before the deadline subject to an energy constraint. Our work in this paper assumes that network parameters such as M, N, λ etc., are known; it will be important to address the adaptive control problem when these parameters are unknown.

APPENDIX A PROOF OF THEOREM 3.1

We first prove that for the optimal policy it is sufficient to consider two actions 1 (i.e., copy now) and *stop* (i.e., do not copy now and never copy again). More precisely, under the optimal policy, if a susceptible relay that is met is not copied, then no susceptible relay is copied in the future as well. Let us fix a $N_0 \leq n \leq N - 1$. Let m_n^* be the maximum j such that $u^*(j, n, r) = 1$.¹⁰ We show that $u^*(j, n, r) = 1$ for all $0 \leq j < m_n^*$; see Figure 2 for an illustration of this fact. The proof is via induction.

Proposition A.1: If $u^*(j, n, r) = 1$ for all $m+1 \leq j \leq m_n^*$, then $u^*(m, n, r) = 1$.

Proof: Define

$$\begin{aligned} \psi(m, n) &:= J_{0s}(m, n, r) - J(m, n, r), \\ \theta_0(m, n) &:= J_{0s}(m, n, r) - A((m, n, r), 0), \\ \text{and } \theta_1(m, n) &:= J_{1s}(m, n, r) - A((m, n, r), 1). \end{aligned}$$

The action sequences that give rise to $J_{0s}(m, n, r)$ and $A((m, n, r), 0)$, do not copy to the susceptible relay that was just met. More formally,¹¹

$$\begin{aligned} J_{0s}(m, n, r) &= C_d((m, n, r), 0) + \sum_{j=m}^{M_\alpha-1} \left(\prod_{l=m}^{j-1} p_{l,n}(d) \right) \\ &\quad \times p_{j,n}(r) \left(\sum_{l=m}^{j-1} C_d((l, n, d), 1) + J_{0s}(j, n, r) \right) \end{aligned}$$

and

¹⁰Note that, for a given n , m_n^* could be 0, in that case we do not copy to any more relays.

¹¹Convention: A sum over an empty index set is 0 and a product over an empty index set is 1, which happen when $j = m$.

$$A((m, n, r), 0) = C_d((m, n, r), 0) + \sum_{j=m}^{M_\alpha-1} \left(\prod_{l=m}^{j-1} p_{l,n}(d) \right) \\ \times p_{j,n}(r) \left(\sum_{l=m}^{j-1} C_d((l, n, d), 1) + J(j, n, r) \right)$$

Thus, subtracting the latter from the former,

$$\theta_0(m, n) = \sum_{j=m}^{M_\alpha-1} \left(\prod_{l=m}^{j-1} p_{l,n}(d) \right) p_{j,n}(r) \psi(j, n). \quad (14)$$

Since $A((m, n, r), 0) \geq J(m, n, r)$, it follows that $\psi(m, n) \geq \theta_0(m, n)$, and so

$$\psi(m, n) \geq \sum_{j=m}^{M_\alpha-1} \left(\prod_{l=m}^{j-1} p_{l,n}(d) \right) p_{j,n}(r) \psi(j, n) \\ = p_{m,n}(r) \psi(m, n) \\ + p_{m,n}(d) \sum_{j=m+1}^{M_\alpha-1} \left(\prod_{l=m+1}^{j-1} p_{l,n}(d) \right) p_{j,n}(r) \psi(j, n)$$

which implies upon rearrangement

$$\psi(m, n) \geq \sum_{j=m+1}^{M_\alpha-1} \left(\prod_{l=m+1}^{j-1} p_{l,n}(d) \right) p_{j,n}(r) \psi(j, n) \quad (15)$$

Next, we establish the following lemma.

Lemma A.1: $\theta_1(m, n) \geq \theta_1(m+1, n)$.

Proof: Note that both the action sequences that lead to the two cost terms in the definition of $\theta_1(m, n)$ copy at state (m, n, r) . Subsequently, both incur equal costs until a decision epoch when an infected node meets a susceptible relay. Also, at any such state $(j, n+1, r)$, $j \geq m$, the costs to go differ by $\psi(j, n+1)$. Hence,

$$\theta_1(m, n) = \sum_{j=m}^{M_\alpha-1} \left(\prod_{l=m}^{j-1} p_{l,n+1}(d) \right) p_{j,n+1}(r) \psi(j, n+1) \\ = p_{m,n+1}(r) \psi(m, n+1) + p_{m,n+1}(d) \theta_1(m+1, n)$$

where

$$\theta_1(m+1, n) = \sum_{j=m+1}^{M_\alpha-1} \left(\prod_{l=m+1}^{j-1} p_{l,n+1}(d) \right) p_{j,n+1}(r) \psi(j, n+1).$$

Thus it suffices to show that

$$\psi(m, n+1) \geq \theta_1(m+1, n),$$

which is same as (15) with n replaced by $n+1$. ■

Next, observe that for all $m \leq j \leq m_n^*$,

$$\psi(j, n) = J_{0s}(j, n, r) - \min\{A((j, n, r), 0), A((j, n, r), 1)\} \\ = \max\{\theta_0(j, n), \Phi(j, n) + \theta_1(j, n)\}. \quad (16)$$

Moreover, from the induction hypothesis, the optimal policy copies at states (j, n, r) for all $m+1 \leq j \leq m_n^*$. Hence, for $m+1 \leq j \leq m_n^*$,

$$\psi(j, n) = \Phi(j, n) + \theta_1(j, n).$$

Finally, $\psi(j, n) = 0$ for all $m_n^* < j \leq M_\alpha - 1$ as the optimal policy does not copy in these states. Hence, from (14),

$$\theta_0(m, n) = p_{m,n}(r) \max\{\theta_0(m, n), \Phi(m, n) + \theta_1(m, n)\} + p_{m,n}(d) \\ \times \sum_{j=m+1}^{m_n^*} \left(\prod_{l=m+1}^{j-1} p_{l,n}(d) \right) p_{j,n}(r) (\Phi(j, n) + \theta_1(j, n)) \\ < p_{m,n}(r) \max\{\theta_0(m, n), \Phi(m, n) + \theta_1(m, n)\} + p_{m,n}(d) \\ \times (\Phi(m, n) + \theta_1(m, n)) \sum_{j=m+1}^{m_n^*} \left(\prod_{l=m+1}^{j-1} p_{l,n}(d) \right) p_{j,n}(r) \\ \leq p_{m,n}(r) \max\{\theta_0(m, n), \Phi(m, n) + \theta_1(m, n)\} \\ + p_{m,n}(d) (\Phi(m, n) + \theta_1(m, n)) \\ = \max\{p_{m,n}(r) \theta_0(m, n) + p_{m,n}(d) (\Phi(m, n) + \theta_1(m, n)), \\ \Phi(m, n) + \theta_1(m, n)\}, \quad (17)$$

where the first (strict) inequality holds because $\Phi(m, n)$ is strictly decreasing (see (4)) and $\theta_1(m, n)$ is decreasing (see Lemma A.1) in m for fixed n . The second inequality follows because the summation term is a probability which is less than 1. Now suppose that $\theta_0(m, n) \geq \Phi(m, n) + \theta_1(m, n)$. Then

$$\max\{p_{m,n}(r) \theta_0(m, n) + p_{m,n}(d) (\Phi(m, n) + \theta_1(m, n)), \\ \Phi(m, n) + \theta_1(m, n)\} \\ = p_{m,n}(r) \theta_0(m, n) + p_{m,n}(d) (\Phi(m, n) + \theta_1(m, n)) \\ \leq \theta_0(m, n)$$

which contradicts (17). Thus, we conclude that

$$\theta_0(m, n) < \Phi(m, n) + \theta_1(m, n).$$

This further implies that $\psi(m, n) = \Phi(m, n) + \theta_1(m, n)$ (see (16)), and so that $u^*(m, n, r) = 1$. ■

We now return to the proof of Theorem 3.1. We show that the one-step look ahead policy is optimal for the resulting stopping problem. To see this, observe that $\Phi(m, n)$ is decreasing in m for a given n and also decreasing in n for a given m . Thus, if $(m, n, r) \in \mathcal{S}_S$, i.e., $\Phi(m, n) \leq 0$ (see (3)), and the susceptible relay that is met is copied, the next state $(m, n+1, r)$ also belongs to the stopping set \mathcal{S}_S . In other words, \mathcal{S}_S is also an absorbing set [19, Section 3.4]. Consequently, the one-step look ahead policy is optimal.

APPENDIX B PROOF OF THEOREM 4.1

We start with a preliminary result and a few definitions.

Proposition B.1: Let $\alpha < 1$ and $Y_0 > 0$. Let ϕ^K and ϕ be as given in (8) and (10), respectively. Then, the functions $\phi^K(\cdot)$ converge to $\phi(\cdot)$ uniformly, i.e., for every $\nu > 0$, there exists a K_ν such that

$$\sup_{(x,y) \in \Delta^K} |\phi^K(x, y) - \phi(x, y)| < \nu$$

for all $K \geq K_\nu$.

Proof: See [24, Appendix B]. ■

In the following, to facilitate a parsimonious description, we use the notation $z^K(t) = (x^K(t), y^K(t))$, $z(t) = (x(t), y(t))$

and $\mathcal{Z} = [0, X_\alpha] \times [Y_0, Y]$. Let us define, for a $\nu \in \mathbb{R}$,

$$\begin{aligned} \mathcal{S}_\nu &= \{z \in \mathcal{Z} : \phi(z) > \nu\}, \\ \tau_\nu &= \inf\{t \geq 0 : z(t) \notin \mathcal{S}_\nu\}, \end{aligned}$$

and a stopping time

$$\tau_\nu^K = \inf\{t \geq 0 : z^K(t) \notin \mathcal{S}_\nu\},$$

the time when $z^K(t)$ exits the limiting set \mathcal{S}_ν . Observe that

$$\frac{\partial \phi}{\partial x} = -\frac{1}{\Lambda(x+y)^2(X-x)} \leq -\frac{1}{\Lambda(X_\alpha+Y)^2 X}, \quad (18)$$

and $f_1^K(x, y)$ (see (7a)) is positive and bounded away from zero. These imply $\tau_\nu^K < \infty$ with probability 1. Similarly, $\tau_\nu < \infty$. The following assertion is a corollary of Proposition B.1.

Corollary B.1: Let K_ν be as in Proposition B.1. For $K \geq K_\nu$,

$$\begin{aligned} \phi^K(z) &> 0 \text{ for all } z \in \mathcal{S}_\nu, \\ \text{and } \phi^K(z) &\leq 0 \text{ for all } z \notin \mathcal{S}_\nu. \end{aligned}$$

We define the uncontrolled dynamics (i.e., the one in which the susceptible relays are always copied) as a Markov process $\bar{z}^K(t) = (\bar{x}^K(t), \bar{y}^K(t))$, $t \geq 0$ for which $\bar{z}^K(0) = z^K(0)$. Let $\bar{z}(t) = (\bar{x}(t), \bar{y}(t))$, $t \geq 0$ be the corresponding limiting deterministic dynamics. Formally, $\bar{z}(0) = z(0)$, and for $t \geq 0$,

$$\begin{aligned} \frac{d\bar{x}(t)}{dt} &= \Lambda(\bar{x}(t) + \bar{y}(t))(X - \bar{x}(t)), \\ \frac{d\bar{y}(t)}{dt} &= \Lambda(\bar{x}(t) + \bar{y}(t))(Y - \bar{y}(t)). \end{aligned}$$

The quantities on the right-hand side of the above equations are at most Λ , and so $\|\frac{d\bar{z}}{dt}\| \leq \sqrt{2}\Lambda$. Also observe that the processes $\bar{z}^K(t)$ and $\bar{z}(t)$ satisfy the hypotheses of Darling [15] (see Section IV-A), and thus convergence of $\bar{z}^K(t)$ to $\bar{z}(t)$ follows.

We also define a Markov process $\tilde{z}^K(t) = (\tilde{x}^K(t), \tilde{y}^K(t))$, $t \geq \tau_\nu$ for which $\tilde{z}^K(\tau_\nu) = z^K(\tau_\nu)$ and

$$\begin{aligned} \frac{d\mathbb{E}(\tilde{x}^K(t) | (\tilde{x}^K(t), \tilde{y}^K(t)))}{dt} &= \Lambda(\tilde{x}^K(t) + \tilde{y}^K(t))(X - \tilde{x}^K(t)) \\ \frac{d\mathbb{E}(y^K(t) | (x^K(t), y^K(t)))}{dt} &= 0 \end{aligned}$$

In other words, $\tilde{z}^K(t)$ is the process in which relays are not copied after τ_ν . Similarly, we define $\tilde{z}(t) = (\tilde{x}(t), \tilde{y}(t))$, $t \geq \tau_\nu$ as the solution of the corresponding differential equations. In other words, $\tilde{z}(\tau_\nu) = z(\tau_\nu)$, and for $t \geq \tau_\nu$,

$$\begin{aligned} \frac{d\tilde{x}(t)}{dt} &= f_1(\tilde{x}(t), \tilde{y}(t)) := \Lambda(\tilde{x}(t) + \tilde{y}(t))(X - \tilde{x}(t)), \\ \frac{d\tilde{y}(t)}{dt} &= f_2(\tilde{x}(t), \tilde{y}(t)) := 0 \end{aligned}$$

We define

$$\begin{aligned} \tilde{\tau}_{-\nu}^K &= \inf\{t \geq \tau_\nu : \tilde{z}^K(t) \notin \mathcal{S}_{-\nu}\}, \\ \tilde{\tau}_{-\nu} &= \inf\{t \geq \tau_\nu : \tilde{z}(t) \notin \mathcal{S}_{-\nu}\}. \end{aligned}$$

Since

$$\Lambda Y_0 (X - X_\alpha) \leq \frac{d\tilde{x}}{dt} \leq \Lambda,$$

TABLE I
VARIABLES AND THEIR DESCRIPTION

variables	description
$z^K(t)$	controlled dynamics with discontinuity at τ^K
$z(t)$	$z^K(t)$'s fluid limit with discontinuity at τ^*
τ_ν^K	instant when $z^K(t)$ exits \mathcal{S}_ν
τ_ν	instant when $z(t)$ exits \mathcal{S}_ν
$\bar{z}^K(t)$	uncontrolled dynamics with no discontinuity
$\bar{z}(t)$	$\bar{z}^K(t)$'s fluid limit with no discontinuity
$\tilde{z}^K(t)$	identical to $z^K(t)$ until τ_ν at which copying to relays is stopped
$\tilde{z}(t)$	$\tilde{z}^K(t)$'s fluid limit with discontinuity at τ_ν
$\tilde{\tau}_{-\nu}^K$	instant when $\tilde{z}^K(t)$ exits $\mathcal{S}_{-\nu}$
$\tilde{\tau}_{-\nu}$	instant when $\tilde{z}(t)$ exits $\mathcal{S}_{-\nu}$

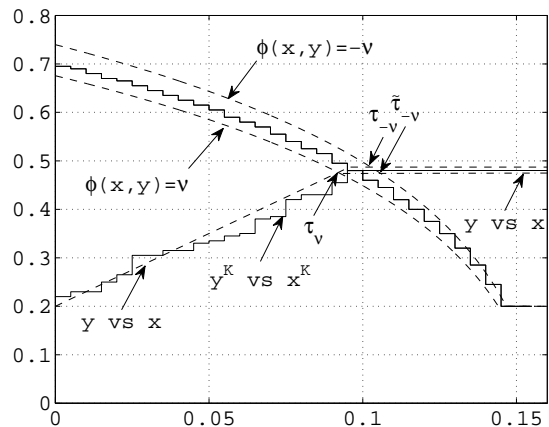


Fig. 10. An illustration of the trajectories of the controlled CTMC and the corresponding ODE, and the associated variables.

the lower bound implies that there is a strictly positive increase in \tilde{x} after time τ_ν . Since $\Phi(x, y)$ decreases with increasing x at a rate bounded away from 0 (see 18), $\tilde{z}(t)$ must exit $\mathcal{S}_{-\nu}$ within a short additional duration. Thus, we have that $\tilde{\tau}_{-\nu} - \tau_\nu \leq b\nu$ for a suitably chosen $b < \infty$.

We summarize the variables in Table I. We also illustrate sample trajectories of a controlled CTMC and the corresponding ODE via an example (Figure 10). We choose $M = 40$, $N = 160$, $\alpha = 0.8$, $N_0 = 40$, $\lambda = 0.00025$ and $\gamma = 0.25$. We plot the graphs of $\phi(x, y) = \nu$ and $\phi(x, y) = -\nu$ for $\nu = 0.2$. We also show the trajectories “ y^K vs x^K ”, “ y vs x ”, “ \tilde{y} vs \tilde{x} ” and the epochs τ_ν , $\tau_{-\nu}$ and $\tilde{\tau}_{-\nu}$.

We also need the following lemmas.

Lemma B.1: For every $\epsilon > 0$, there exists a $\bar{\tau}_\epsilon$ such that for all $t \geq 0$, $0 \leq u \leq \bar{\tau}_\epsilon$,

$$\mathbb{P}(\|\bar{z}^K(t+u) - \bar{z}(t)\| > \epsilon) = O(K^{-1}).$$

Proof: See [24, Appendix B]. ■

Lemma B.2: Suppose u is a fixed time and u^K is a random time that satisfies $\mathbb{P}(|u - u^K| > \delta) = O(K^{-1})$ for every $\delta >$

0. Then, for every $\epsilon > 0$,

$$\mathbb{P}(\|\bar{z}^K(u) - \bar{z}^K(u \wedge u^K)\| > \epsilon) = O(K^{-1})$$

Proof: See [24, Appendix B]. \blacksquare

We now prove the assertion in Theorem 4.1 in three steps:

(a) over $[0, \tau_\nu]$, (b) over $[\tau_\nu, \tilde{\tau}_{-\nu}]$ and (c) over $[\tilde{\tau}_{-\nu}, \tau]$.

(a) First, we prove the convergence of $z^K(t)$ to $z(t)$ over $[0, \tau_\nu]$. Fix a $\nu > 0$. Then Corollary B.1 implies that $z^K(t)$ converges to $z(t)$ in the region \mathcal{S}_ν . Following [15, Theorem 2.8] we have, for all $\epsilon, \delta > 0$,

$$\mathbb{P}\left(\sup_{0 \leq t \leq \tau_\nu} \|z^K(t \wedge \tau_\nu^K) - z(t)\| > \epsilon\right) = O(K^{-1}) \quad (19a)$$

$$\text{and } \mathbb{P}(|\tau_\nu^K - \tau_\nu| > \delta) = O(K^{-1}). \quad (19b)$$

Since, for all $t \geq 0$,

$$\|z^K(t) - z(t)\| \leq \|z^K(t \wedge \tau_\nu^K) - z(t)\| + \|z^K(t) - z^K(t \wedge \tau_\nu^K)\|,$$

we obtain

$$\sup_{0 \leq t \leq \tau_\nu} \|z^K(t) - z(t)\| \leq \sup_{0 \leq t \leq \tau_\nu} \|z^K(t \wedge \tau_\nu^K) - z(t)\| + \sup_{0 \leq t \leq \tau_\nu} \|z^K(t) - z^K(t \wedge \tau_\nu^K)\|.$$

If the left side is larger than ϵ , at least one of the two terms on the right side is larger than $\epsilon/2$, and so by the union bound, we get

$$\begin{aligned} & \mathbb{P}\left(\sup_{0 \leq t \leq \tau_\nu} \|z^K(t) - z(t)\| > \epsilon\right) \\ & \leq \mathbb{P}\left(\sup_{0 \leq t \leq \tau_\nu} \|z^K(t \wedge \tau_\nu^K) - z(t)\| > \frac{\epsilon}{2}\right) \\ & \quad + \mathbb{P}\left(\sup_{0 \leq t \leq \tau_\nu} \|z^K(t) - z^K(t \wedge \tau_\nu^K)\| > \frac{\epsilon}{2}\right) \\ & \leq O(K^{-1}) + \mathbb{P}\left(\|z^K(\tau_\nu) - z^K(\tau_\nu \wedge \tau_\nu^K)\| > \frac{\epsilon}{2}\right) \end{aligned} \quad (20)$$

where the first term in the last inequality follows from (19a). Also, from corollary B.1, for $K \geq K_\nu$, $\phi^K(z^K(\tau_\nu^K)) > 0$, i.e., the process $z^K(t)$ follows uncontrolled dynamics until τ_ν^K . Thus, for $K \geq K_\nu$, $z^K(\tau_\nu^K) = \bar{z}^K(\tau_\nu^K)$ and

$$\|z^K(\tau_\nu) - z^K(\tau_\nu \wedge \tau_\nu^K)\| \leq \|\bar{z}^K(\tau_\nu) - \bar{z}^K(\tau_\nu \wedge \tau_\nu^K)\|$$

sample path wise. The inequality is an equality if $\tau_\nu \leq \tau_\nu^K$; both sides equal 0 in this case. Otherwise, it is an inequality because the possible change in dynamics of $z^K(t)$ after τ_ν^K makes it increase (in both its components) at a slower pace than the uncontrolled $\bar{z}^K(t)$. Thus

$$\begin{aligned} & \mathbb{P}\left(\|z^K(\tau_\nu) - z^K(\tau_\nu \wedge \tau_\nu^K)\| > \frac{\epsilon}{2}\right) \\ & \leq \mathbb{P}\left(\|\bar{z}^K(\tau_\nu) - \bar{z}^K(\tau_\nu \wedge \tau_\nu^K)\| > \frac{\epsilon}{2}\right) \leq O(K^{-1}), \end{aligned}$$

where the last inequality follows from (19b) and Lemma B.2. Using this in (20) we get

$$\begin{aligned} \mathbb{P}\left(\sup_{0 \leq t \leq \tau_\nu} \|z^K(t) - z(t)\| > \epsilon\right) & \leq O(K^{-1}) + O(K^{-1}) \\ & = O(K^{-1}) \end{aligned}$$

(b) Now we prove the convergence of $z^K(t)$ to $z(t)$ over $[\tau_\nu, \tilde{\tau}_{-\nu}]$. Observe that, for $t \in [\tau_\nu, \tilde{\tau}_{-\nu}]$,

$$\begin{aligned} & \|z^K(t) - z(t)\| \\ & \leq \|z^K(\tau_\nu) - z(\tau_\nu)\| + \|z^K(t) - z^K(\tau_\nu)\| + \|z(t) - z(\tau_\nu)\|. \end{aligned}$$

Hence,

$$\begin{aligned} & \sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z^K(t) - z(t)\| \\ & \leq \|z^K(\tau_\nu) - z(\tau_\nu)\| + \sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z^K(t) - z^K(\tau_\nu)\| \\ & \quad + \sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z(t) - z(\tau_\nu)\| \\ & = \|z^K(\tau_\nu) - z(\tau_\nu)\| + \|z^K(\tilde{\tau}_{-\nu}) - z^K(\tau_\nu)\| \\ & \quad + \|z(\tilde{\tau}_{-\nu}) - z(\tau_\nu)\| \\ & \leq \|z^K(\tau_\nu) - z(\tau_\nu)\| + \|z^K(\tilde{\tau}_{-\nu}) - z^K(\tau_\nu)\| + \sqrt{2}\Lambda b\nu \end{aligned}$$

where the equality follows because the $z(t)$ and $z^K(t)$ are nondecreasing. The last inequality holds because $\|dz/dt\| \leq \|d\bar{z}/dt\| \leq \sqrt{2}\Lambda$ and $\tilde{\tau}_{-\nu} - \tau_\nu \leq b\nu$. Moreover,

$$\begin{aligned} & \mathbb{P}\left(\sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z^K(t) - z(t)\| > \sqrt{2}\Lambda b\nu + \frac{\epsilon}{2}\right) \\ & \leq \mathbb{P}\left(\|z^K(\tau_\nu) - z(\tau_\nu)\| > \frac{\epsilon}{4}\right) \\ & \quad + \mathbb{P}\left(\|z^K(\tilde{\tau}_{-\nu}) - z^K(\tau_\nu)\| > \frac{\epsilon}{4}\right) \\ & = O(K^{-1}) + \mathbb{P}\left(\|z^K(\tilde{\tau}_{-\nu}) - z^K(\tau_\nu)\| > \frac{\epsilon}{4}\right) \end{aligned}$$

where the equality follows from the result of part (a). We now redefine the Markov process $\bar{z}^K(t) = (\bar{x}^K(t), \bar{y}^K(t))$ for $t \geq \tau_\nu$, to be the uncontrolled dynamics with initial condition $\bar{z}^K(\tau_\nu) = z^K(\tau_\nu)$. Again, it can be easily observed that

$$\|z^K(\tilde{\tau}_{-\nu}) - z^K(\tau_\nu)\| \leq \|\bar{z}^K(\tilde{\tau}_{-\nu}) - \bar{z}^K(\tau_\nu)\|.$$

Thus

$$\begin{aligned} & \mathbb{P}\left(\sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z^K(t) - z(t)\| > \sqrt{2}\Lambda b\nu + \frac{\epsilon}{2}\right) \\ & \leq O(K^{-1}) + \mathbb{P}\left(\|\bar{z}^K(\tilde{\tau}_{-\nu}) - \bar{z}^K(\tau_\nu)\| > \frac{\epsilon}{4}\right) \\ & \leq O(K^{-1}) + \mathbb{P}\left(\|\bar{z}^K(\tau_\nu + b\nu) - \bar{z}^K(\tau_\nu)\| > \frac{\epsilon}{4}\right) \end{aligned}$$

Set $\nu = \min\{\frac{\epsilon}{2\sqrt{2}\Lambda b}, \frac{\tau_{\frac{\epsilon}{4b}}}{4b}\}$, and apply Lemma B.1 to get

$$\begin{aligned} & \mathbb{P}\left(\sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z^K(t) - z(t)\| > \epsilon\right) \\ & \leq \mathbb{P}\left(\sup_{\tau_\nu \leq t \leq \tilde{\tau}_{-\nu}} \|z^K(t) - z(t)\| > \sqrt{2}\Lambda b\nu + \frac{\epsilon}{2}\right) \\ & \leq O(K^{-1}) + O(K^{-1}) = O(K^{-1}). \end{aligned}$$

(c) Finally, we prove the convergence of $z^K(t)$ to $z(t)$ over $[\tilde{\tau}_{-\nu}, \tau]$. Reconsider the process $\bar{z}^K(t)$, $t \geq \tau_\nu$ and the associated function $\tilde{z}(t)$. Recall that, for any $\nu > 0$, $\bar{z}^K(t)$ and $\tilde{z}(t)$ exit $\mathcal{S}_{-\nu}$ at $\tilde{\tau}_{-\nu}^K$ and $\tilde{\tau}_{-\nu}$ respectively. Clearly, $\tilde{\tau}_{-\nu/2} < \tilde{\tau}_{-\nu}$; say $\tilde{\tau}_{-\nu} - \tilde{\tau}_{-\nu/2} = \delta_\nu$. Also, using [15,

Theorem 2.8],

$$\mathbb{P}\left(\tilde{\tau}_{-\nu/2}^K - \tilde{\tau}_{-\nu/2} > \delta_\nu\right) = O(K^{-1})$$

$$\text{i.e., } \mathbb{P}\left(\tilde{\tau}_{-\nu/2}^K > \tilde{\tau}_{-\nu}\right) = O(K^{-1})$$

Furthermore, $\tau_{-\nu/2}^K \leq \tilde{\tau}_{-\nu/2}^K$ sample path wise. The inequality holds because $z^K(t)$ may continue to increase (in both its components) at a higher pace than $\tilde{z}^K(t)$ even after τ_ν . Thus

$$\mathbb{P}\left(\tau_{-\nu/2}^K > \tilde{\tau}_{-\nu}\right) = O(K^{-1}),$$

implying that the probability that $z^K(t)$ has changed its dynamics by $\tilde{\tau}_{-\nu}$ approaches 1 as K approaches ∞ . In these realizations, the dynamics of $z^K(t)$ and $z(t)$ match for $t \geq \tilde{\tau}_{-\nu}$. We restrict ourselves to only these realizations. We also have from part (b) that, for every $\epsilon > 0$,

$$\mathbb{P}\left(\|z^K(\tilde{\tau}_{-\nu}) - z(\tilde{\tau}_{-\nu})\| > \epsilon\right) = O(K^{-1})$$

Once more using [15, Theorem 2.8], for any $\epsilon, \delta > 0$

$$\mathbb{P}\left(\sup_{\tilde{\tau}_{-\nu} \leq t \leq \tau} \|z^K(t) - z(t)\| > \epsilon\right) = O(K^{-1})$$

and $\mathbb{P}\left(|\tau^K - \tau| > \delta\right) = O(K^{-1})$.

APPENDIX C

PROOF OF THEOREM 4.2

Observe that $\mathcal{T}_d = \tau^K$ by definition (see (11)), and that all the destinations are copied under any policy. Hence, the total expected cost under the optimal policy u^* is

$$\mathbb{E}_{u^*}^K\{\mathcal{T}_d + \gamma\mathcal{E}_c\} = \mathbb{E}_{u^*}^K\{\tau^K + \Gamma y^K(\tau^K)\}.$$

Under the deterministic policy u^∞ , copying to relays is stopped at the deterministic time instant τ^* . So, it incurs the total expected cost

$$\mathbb{E}_{u^\infty}^K\{\mathcal{T}_d + \gamma\mathcal{E}_c\} = \mathbb{E}_{u^\infty}^K\{\tau^K + \Gamma y^K(\tau^*)\}.$$

Also, under u^∞ , the fluid limits of $(x^K(t), y^K(t))$ are the same deterministic dynamics $(x(t), y(t))$ defined in Section IV-A (i.e., solutions of (9a)-(9b)). Indeed, $(x^K(t), y^K(t))$ and $(x(t), y(t))$ satisfy the hypotheses assumed in Darling [15] over the intervals $[0, \tau^*]$ and $[\tau^*, \infty)$. Thus [15, Theorem 2.8] applies, and we obtain¹²

$$\lim_{K \rightarrow \infty} \mathbb{P}_{u^\infty}^K\left(\sup_{0 \leq t \leq \tau} \|(x^K(t), y^K(t)) - (x(t), y(t))\| > \epsilon\right) = 0,$$

$$\lim_{K \rightarrow \infty} \mathbb{P}_{u^\infty}^K\left(|\tau^K - \tau| > \delta\right) = 0.$$

We first show that, under the control u^* , $y^K(\tau^K)$ converge to $y(\tau^*)$ in probability. Recall that $\phi(x(t), y(t))$ is decreasing in t , $\phi(x(\tau^*), y(\tau^*)) = 0$ (see (13)) and $\phi(x(\tau), y(\tau)) = -\Gamma$ (see (10) and (12)). Thus $\tau^* < \tau$, and from (9b), $y(\tau) = y(\tau^*)$. Hence it suffices to show that $y^K(\tau^K)$ converge to $y(\tau)$ in probability. To see this, observe that

$$|y^K(\tau^K) - y(\tau)| \leq |y^K(\tau^K) - y^K(\tau)| + |y^K(\tau) - y(\tau)|. \quad (21)$$

¹²Applying [15, Theorem 2.8] over $[0, \tau^*]$ yields $\lim_{K \rightarrow \infty} \mathbb{P}\left(\|(x^K(\tau^*), y^K(\tau^*)) - (x(\tau^*), y(\tau^*))\| > \epsilon\right) = 0$ which is a necessary condition to apply [15, Theorem 2.8] over $[\tau^*, \infty)$.

From Theorem 4.1, $y^K(\tau)$ and τ^K converge to $y(\tau)$ and τ respectively, in probability. The latter result and arguments similar to those in the proof of Lemma B.2 imply that

$$\mathbb{P}\left(|y^K(\tau^K) - y^K(\tau)| > \epsilon\right) = O(K^{-1})$$

for every $\epsilon > 0$. Using these facts in (21), we conclude that

$$\mathbb{P}\left(|y^K(\tau^K) - y(\tau)| > \epsilon\right) = O(K^{-1}).$$

for every $\epsilon > 0$ which is the desired claim.

Further, $y^K(\tau^K)$ are bounded uniformly over all K . Thus, following [26, Remark 9.5.1], $y^K(\tau^K)$ are uniformly integrable under u^* . Similar arguments imply that, under u^∞ also, $y^K(\tau^*)$ converge in probability to $y(\tau^*)$, and are uniformly integrable. Then, the convergence in probability along with [26, Theorem 9.5.1] implies that

$$\lim_{K \rightarrow \infty} \mathbb{E}_{u^*}^K y^K(\tau^K) = \lim_{K \rightarrow \infty} \mathbb{E}_{u^\infty}^K y^K(\tau^*) = y(\tau^*). \quad (22)$$

Next, it is shown that under both the controls u^* and u^∞ , the delivery delays τ^K have second moments that are bounded uniformly over all K . To see this, consider a policy u^0 that never copies to relays. Clearly,

$$\mathbb{E}_{u^*}^K(\tau^K)^2 < \mathbb{E}_{u^0}^K(\tau^K)^2,$$

$$\mathbb{E}_{u^\infty}^K(\tau^K)^2 < \mathbb{E}_{u^0}^K(\tau^K)^2$$

for each K . Thus it suffices to show that

$$\sup_K \mathbb{E}_{u^0}^K(\tau^K)^2 < \infty. \quad (23)$$

Note that

$$\tau^K = \sum_{m=0}^{M_\alpha^K - 1} \bar{\delta}_m,$$

where $\bar{\delta}_m$ is the time duration for which $m(t) = m$. Under the policy u^0 , $\bar{\delta}_m$ is exponentially distributed with mean $\frac{1}{\lambda^K(m + N_0^K)(M^K - m)}$. So,

$$\begin{aligned} \mathbb{E}_{u^0}^K \tau^K &= \sum_{m=0}^{M_\alpha^K - 1} \frac{1}{\lambda^K(m + N_0^K)(M^K - m)} \\ &\leq \sum_{m=0}^{M_\alpha^K - 1} \frac{1}{\lambda^K N_0^K (M^K - M_\alpha^K)} \\ &= \frac{M_\alpha^K}{\lambda^K N_0^K (M^K - M_\alpha^K)} \\ &= \frac{X_\alpha}{\Lambda Y_0 (X - X_\alpha)} < \infty. \end{aligned}$$

Also, $\bar{\delta}_m, m = 0, 1, \dots$ are independent. Thus,

$$\begin{aligned} \text{Var}_{u^0}^K \tau^K &= \sum_{m=0}^{M_\alpha^K - 1} \frac{1}{(\lambda^K(m + N_0^K)(M^K - m))^2} \\ &\leq \frac{M_\alpha^K}{(\lambda^K N_0^K (M^K - M_\alpha^K))^2} \\ &= \frac{X_\alpha}{K \Lambda^2 Y_0^2 (X - X_\alpha)^2} \rightarrow 0 \end{aligned}$$

as $K \rightarrow \infty$. These results together imply (23).

Again, from [26, Remark 9.5.1], τ^K are uniformly integrable under both u^* and u^∞ . Recall from Theorem 4.1 that τ^K converge to τ in probability. Once more using [26, Theorem 9.5.1],

$$\lim_{K \rightarrow \infty} \mathbb{E}_{u^*}^K \tau^K = \lim_{K \rightarrow \infty} \mathbb{E}_{u^\infty}^K \tau^K = \tau. \quad (24)$$

Finally, combining (22) and (24), we conclude

$$\lim_{K \rightarrow \infty} \mathbb{E}_{u^*}^K \{\mathcal{T}_d + \gamma \mathcal{E}_c\} = \lim_{K \rightarrow \infty} \mathbb{E}_{u^\infty}^K \{\mathcal{T}_d + \gamma \mathcal{E}_c\} = \tau + \Gamma y(\tau^*).$$

REFERENCES

- [1] K. Fall, "A delay-tolerant network architecture for challenged internets," in *SIGCOMM 03*, Karlsruhe, Germany, 2003, pp. 27–34.
- [2] A. Vahdat and D. Becker, "Epidemic routing for partially-connected ad hoc networks," Duke University, Tech. Rep. CS-2000-06, July 2000.
- [3] S. Venkatramanan and A. Kumar, "Co-evolution of content popularity and delivery in mobile P2P networks," in *Proceedings of IEEE INFOCOM*, Orlando, FL, USA, March 2012.
- [4] A. Karnik and P. Dayama, "Optimal control of information epidemics," in *Proceedings of COMSNETS*, Bangalore, India, January 2012.
- [5] M. Khouzani, S. Sarkar, and E. Altman, "Dispatch then stop: Optimal dissemination of security patches in mobile wireless networks," in *Proceedings of IEEE CDC*, Atlanta, Georgia, USA, December 2010.
- [6] R. Groenevelt, P. Nain, and G. Koole, "The message delay in mobile ad hoc networks," *Performance Evaluation*, vol. 62, no. 1–4, pp. 210–228, October 2005.
- [7] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," *Computer Networks*, vol. 51, no. 10, pp. 2867–2891, July 2007.
- [8] G. Neglia and X. Zhang, "Optimal delay-power tradeoff in sparse delay tolerant networks: a preliminary study," in *SIGCOMM Workshop on Challenged Networks*, Pisa, Italy, September 2006, pp. 237–244.
- [9] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: The multiple-copy case," *IEEE/ACM Transactions on Networking*, vol. 16, no. 1, pp. 77–90, February 2008.
- [10] E. Altman, T. Basar, and F. D. Pellegrini, "Optimal monotone forwarding policies in delay tolerant mobile ad-hoc networks," *Performance Evaluation*, vol. 67, no. 4, pp. 299–317, April 2010.
- [11] E. Altman, G. Neglia, F. D. Pellegrini, and D. Miorandi, "Decentralized stochastic control of delay tolerant networks," in *Proceedings of IEEE INFOCOM*, Rio de Janeiro, Brazil, April 2009.
- [12] E. Altman, A. P. Azad, F. D. Pellegrini, and T. Basar, "Optimal activation and transmission control in delay tolerant networks," in *Proceedings of IEEE INFOCOM*, San Diego, CA, USA, March 2010.
- [13] Y. Li, Y. Jiang, D. Jin, L. Su, L. Zeng, and D. Wu, "Energy-efficient optimal opportunistic forwarding for delay-tolerant networks," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 9, pp. 4500–4512, November 2010.
- [14] T. G. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," *Journal of Applied Probability*, vol. 7, pp. 49–58, 1970.
- [15] R. W. R. Darling, "Fluid limits of pure jump Markov processes: a practical guide," arXiv:math/0210109v3, 2002.
- [16] R. W. R. Darling and J. R. Norris, "Differential equation approximations for Markov chains," *Probability Surveys*, vol. 5, pp. 37–79, 2008.
- [17] N. Gast and B. Gaujal, "Mean field limit of non-smooth systems and differential inclusions," INRIA, Tech. Rep., 2010.
- [18] N. Gast, B. Gaujal, and J.-Y. L. Boudec, "Mean field for markov decision processes: from discrete to continuous optimization," *IEEE Transactions on Automatic Control*, vol. 57, no. 9, pp. 2266–2280, September 2012.
- [19] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. 2*, 3rd ed. Belmont, Massachusetts: Athena Scientific, 2007.
- [20] C. Singh, A. Kumar, and R. Sundaresan, "Delay and energy optimal two-hop relaying in delay tolerant networks," in *8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, Avignon, France, May–June 2010, pp. 256–265.
- [21] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad-hoc wireless networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, pp. 477–486, August 2002.
- [22] T. Karagiannis, J.-Y. L. Boudec, and M. Vojnovic, "Power law and exponential decay of intercontact times between mobile devices," *IEEE Transactions on Mobile Computing*, vol. 9, no. 10, pp. 1377–1390, October 2010.
- [23] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. 1*, 3rd ed. Belmont, Massachusetts: Athena Scientific, 2005.
- [24] C. Singh, E. Altman, A. Kumar, and R. Sundaresan, "Optimal forwarding in delay tolerant networks with multiple destinations," Indian Institute of Science, Bangalore, India, Tech. Rep., 2012, <http://arxiv.org/abs/1110.6127>.
- [25] R. Solis, V. S. Borkar, and P. R. Kumar, "A new distributed time synchronization protocol for multihop wireless networks," in *Proceedings of IEEE CDC*, San Diego, CA, USA, December 2006.
- [26] K. B. Athreya and S. N. Lahiri, *Measure Theory and Probability Theory*. Springer, 2006.