

Capacity and Delay Tradeoffs for Ad-Hoc Mobile Networks

Michael J. Neely, *Member, IEEE*, and Eytan Modiano, *Senior Member, IEEE*

Abstract—We consider the throughput/delay tradeoffs for scheduling data transmissions in a mobile ad-hoc network. To reduce delays in the network, each user sends redundant packets along multiple paths to the destination. Assuming the network has a cell partitioned structure and users move according to a simplified independent and identically distributed (i.i.d.) mobility model, we compute the exact network capacity and the exact end-to-end queueing delay when no redundancy is used. The capacity achieving algorithm is a modified version of the Grossglauser-Tse 2-hop relay algorithm and provides $O(N)$ delay (where N is the number of users). We then show that redundancy cannot increase capacity, but can significantly improve delay. The following necessary tradeoff is established: $\text{delay}/\text{rate} \geq O(N)$. Two protocols that use redundancy and operate near the boundary of this curve are developed, with delays of $O(\sqrt{N})$ and $O(\log(N))$, respectively. Networks with non-i.i.d. mobility are also considered and shown through simulation to closely match the performance of i.i.d. systems in the $O(\sqrt{N})$ delay regime.

Index Terms—fundamental limits, queueing analysis, stochastic systems, wireless networks

I. INTRODUCTION

In this paper we consider the effects of transmitting redundant packets through multiple paths of an ad-hoc wireless network with mobility. Such redundancy improves delay at the expense of increasing overall network congestion. We show that redundancy cannot increase network capacity, but can significantly improve delay performance, yielding delay reductions by several orders of magnitude when data rates are sufficiently less than capacity.

We use the following *cell partitioned* network model: The network is partitioned into C non-overlapping cells of equal size (see Fig. 1). There are N mobile users independently roaming from cell to cell over the network, and time is slotted so that users remain in their current cells for a timeslot, and potentially move to a new cell at the end of the slot. If two users are within the same cell during a timeslot, one can




Fig. 1. A cell-partitioned ad-hoc wireless network with C cells and N mobile users.

transfer a single packet to the other. Each cell can support exactly one packet transfer per timeslot, and users within different cells cannot communicate during the slot. Multi-hop packet transfer proceeds as users change cells and exchange data. The cell partitioning reduces scheduling complexity and facilitates analysis. Similar cell partitioning has recently been considered by Cruz et. al in [4].

We consider the following simplified mobility model: Every timeslot, users choose a new cell location independently and identically distributed over all cells in the network. Such a mobility model is, of course, an over-simplification. Indeed, actual mobility is better described by Markovian dynamics, where users choose new locations every timeslot from the set of cells adjacent to their current cell. However, analysis under the simplified *independent and identically distributed* (i.i.d.) mobility model provides a meaningful bound on performance in the limit of *infinite mobility*. With this assumption, the network topology dramatically changes every timeslot, so that network behavior cannot be predicted and fixed routing algorithms cannot be used. Rather, because information about the current and future locations of users is unknown, one must rely on robust scheduling algorithms. Furthermore, it is shown in [1] [5] that network capacity depends only on the steady-state channel distribution, and hence the capacity region under an i.i.d. mobility model is identical to the capacity of a network with non-i.i.d. mobility with the same steady state distribution (see Corollary 5 on page 88 of [1]). Thus, our capacity results hold also for cases where mobility is described by simple Markovian random walks, considered in Sections VII and VIII. Delay analysis for non-i.i.d. mobility is also presented, and simulations demonstrate that throughput and delay performance is qualitatively similar to the i.i.d. case.

We compute an exact expression for the per-user transmission capacity of the network (for any number of users $N \geq 3$), and show that this capacity cannot be increased by using redundant packet transfers. When no redundancy is

Manuscript received November 19, 2003; revised November 9, 2004. This work was supported in part by NSF ITR grant CCR-0325401 and by DARPA/AFOSR through the University of Illinois grant no. F49620-02-1-0325 for the project entitled "Cooperative Networked Control of Dynamical Peer-to-Peer Vehicle Systems." This work was presented in part at the Conference on Information Sciences and Systems, Johns Hopkins University, March 2003, and as an invited paper at the IEEE Broadnets Conference, San Jose, CA, October 2004.

M. J. Neely is with the Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 USA (email: mjneely@usc.edu, web: <http://www-ref.usc.edu/~mjneely>).

E. Modiano is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (email: modiano@mit.edu).

Communicated by G. Sasaki, Associate Editor for Communication Networks.

Due to a compiling error, the original printed version of this paper incorrectly shifted all bibliographic references after reference [13] by one. The problem has been fixed in this electronic version of the paper.

used, a modified version of the Grossglauser-Tse 2-hop relay algorithm in [6] is presented and shown to achieve capacity. The queueing delay in the network is explicitly computed and shown to be $O(N)/(\mu - \lambda_i)$ (where μ is the per-user network capacity, and λ_i is the rate at which user i transfers packets intended for its destination). Furthermore, it is shown that no scheduling algorithm can improve upon $O(N)$ delay performance unless redundancy is used.

We then consider modifying the 2-hop relay algorithm to allow redundant packet transmissions. It is shown that no scheme which restricts packets to two hops can achieve a better delay than $O(\sqrt{N})$. A scheduling protocol that employs redundant packet transmissions is developed and shown to achieve this delay bound when all users communicate at a reduced data rate of $O(1/\sqrt{N})$. A multi-hop protocol is then developed to achieve $O(\log(N))$ delay by further sacrificing throughput. Finally, the necessary condition $\text{delay}/\text{rate} \geq O(N)$ is established for any routing and scheduling algorithm, and the 2-hop relay algorithms are shown to meet this bound with equality while the multi-hop algorithm deviates from optimality by no more than a logarithmic factor.

Previous work on the capacity of ad-hoc wireless networks is found in [1]–[12], [14]. Gupta and Kumar present asymptotic results for static networks in [7], [8], where it is shown that per-user network capacity is $O(1/\sqrt{N})$, and hence vanishes as the number of users N increases. The effect of mobility on the capacity of ad-hoc wireless networks was first explicitly developed in [6], where a 2-hop relay algorithm was developed and shown to support constant per-user throughput which does not vanish as the size of the network grows. These works do not consider the associated network delay, and analysis of the fundamental queueing delay bounds for general networks remains an important open question.

In [9] it is shown that for a network with a mixture of stationary users and mobile relay nodes, delay can be improved by exploiting velocity information and relaying packets to nodes moving in the direction of their destination. Routing for fully mobile networks using table updates is considered in [10]. Schemes for improving delay via diversity coding and multi-path routing are considered in [11], [12], although these do not consider delays due to path sharing, queueing, or stochastic arrivals. Delay improvement via redundant packet transfers is considered in [3]. This idea is related to the notion of *content replication* considered for static peer-to-peer systems in [13] and for mobile networks in [14]. Our i.i.d. mobility model is similar to that used in [14], where *mobile infostations* are used to store content for users requesting file access. Throughput and delay tradeoffs were perhaps first considered in [15], where delay of multi-hop routing is reduced by increasing the coverage radius of each transmission, at the expense of reducing the number of simultaneous transmissions the network can support. Similar radial scaling techniques have recently appeared in [16] [17] [18]. While our work was developed prior to the work in [16] [17] [18] and does not directly consider radial scaling, for completeness we include a detailed comparison with these approaches at the end of Section VI.

In this paper, we analyze the capacity and delay of cell

partitioned networks and consider the full effects of queueing. Our contributions are threefold: First, we develop an expression for network capacity and compute the exact delay of a capacity achieving strategy. Second, we demonstrate that redundant packet transfers can significantly reduce delay at the cost of reducing throughput. Third, we establish a fundamental delay/rate tradeoff curve that bounds performance of any routing and scheduling algorithm. Protocols for three different rate regimes are developed and shown to operate on or near the boundary of this curve.

A. Concerning the Cell-Partitioned Network Assumptions

The cell partitioned model is used to enable a simple and insightful network analysis, and not necessarily to propose a practical communication scheme. While a direct implementation of a cell partitioned network simplifies scheduling decisions by enabling control actions to be independently distributed over each cell, extra management is required to maintain the cell structure and to coordinate communication between mobile users. First, mobile users must determine their own cell locations. This might be accomplished through satellite positioning signals, assuming each user is equipped with a GPS receiver. Alternatively, location might be determined by triangulating against pre-established ground beacons. These ground beacons could additionally act as *control stations* that handle control signaling between users. We assume that control information for each cell is passed over reserved bandwidth channels.

The cell partitioned network model restricts communication to one transmission per cell per timeslot. This restriction alleviates the interference problems associated with two users simultaneously transmitting in the same cell. However, this does not solve the inter-cell interference problem, as a transmitter in one cell may be very close to a receiver in its neighboring cell. Such interference can be mitigated by requiring users in neighboring cells to transmit over orthogonal frequency bands. It is well known that for rectilinear cell partitionings as in Fig. 1, only four separate frequency bands are needed to ensure that no neighboring cells use the same frequency, and this number can be reduced to three if cells are arranged in a hexagonal pattern. Additional frequency bands can be added to increase the frequency reuse distance, at the cost of reducing the bit rate of each user-to-user transmission. From a theoretical perspective, we note that the capacity expression derived for cell partitioned networks in Section II is very close to the maximum throughput estimates of the Grossglauser-Tse relay strategy, which uses a nearest neighbor transmission policy for networks with full interference and no bandwidth subdivision. Thus, cell partitioned networks can serve as useful theoretical models for analyzing more complex systems, and the protocols we develop for cell partitioned networks can be applied to these other settings.

Throughout this paper, we assume that the number of cells is of the same order as the number of users, so that the user/cell density d is constrained to be $O(1)$ (independent of N). This is a necessary constraint in cases when the network area is increased while maintaining the same average number

of users per unit area and the same transmission power (and hence, transmission radius) for each user. In the opposite case when the network area is fixed but the number of users N grows large (increasing the number of users per unit area), it is possible to consider cell densities that increase with N , although the $d = O(1)$ constraint can still be imposed by appropriately scaling the cell size. Note that in this case, it would be possible to maintain a cell size for which the user/cell density increases to infinity with N . However, this would require the coordination of an increasingly large number of users in each cell, and it would necessarily shrink network capacity to zero with growing N (as shown in the next section). It could, however, provide an alternate means of improving network delay, as described in [16] [17] [18]. Indeed, in the extreme case where there is only one cell containing all nodes, it is clear that any user could reach any other user in just a single hop. A detailed comparison of our results with those of [16] [17] [18] is given in Section VI.

B. Paper Outline

In the next section, we establish the capacity of the cell partitioned network and analyze the delay of the capacity achieving relay algorithm. In Section III we develop delay bounds for transmission schemes with redundancy, and in Section IV we provide scheduling protocols which achieve these bounds. In Section V we prove necessity of $\text{delay}/\text{rate} \geq O(N)$, and show that the given protocols operate on the boundary of this rate-delay tradeoff curve. Simulations and Markovian mobility models are considered in Sections VI and VII.

II. CAPACITY, DELAY, AND THE 2-HOP RELAY ALGORITHM

Consider a cell partitioned network such as that of Fig. 1. The shape and layout of cell regions is arbitrary, although we assume that cells have identical area, do not overlap, and completely cover the network area. We define:

- N = Number of Mobile Users
- C = Number of Cells
- $d = N/C$ = User/Cell density

Users move independently according to the *full-mobility model*, where the steady state location of each user is uniform over all cells.

Let λ_i represent the exogenous arrival rate of packets to user i (in units of packets/slot). Packets are assumed to arrive as a Bernoulli process, so that with probability λ_i a single packet arrives during the current slot, and otherwise no packet arrives. Other stochastic inputs with the same time average arrival rate can be treated similarly, and the arrival model does not affect the region of rates the network can support (see [1]).

We assume packets from source i must be delivered to a unique destination j . In particular we assume the number of users N is even and consider the one-to-one pairing: $1 \leftrightarrow 2$, $3 \leftrightarrow 4$, ..., $(N-1) \leftrightarrow N$; so that user 1 communicates with user 2 and user 2 communicates with user 1, user 3 communicates with user 4 and user 4 communicates with user

3, and so on. Other source-destination scenarios can be treated similarly (see Section II-B).

Packets are transmitted and routed through the network according to some scheduling algorithm. The algorithm chooses which packets to transmit on each timeslot without violating the physical constraints of the cell partitioned network or the following additional *causality constraint*: A user cannot transmit a packet that it has never received. Note that once a packet has been received by a user, it can be stored in memory and transmitted again and again if so desired. We assume that packets are equipped with header information so that they can be individually distinguished for scheduling purposes.

A scheduling algorithm is *stable* if the λ_i rates are satisfied for all users so that queues do not grow to infinity and average delays are bounded. Assuming that all users receive packets at the same data rate (so that $\lambda_i = \lambda$ for all i), the *capacity* of the network is the maximum rate λ that the network can stably support. Note that this is a purely network layer notion of capacity, where optimization is over all possible routing and scheduling protocols. Below we compute the network capacity, assuming users change cells in an i.i.d. fashion every timeslot. In [1] [5] it is shown that the capacity region depends only on the steady state user location distribution. Hence, any Markovian model of user mobility which in steady state distributes users independently and uniformly over the network yields the same expression for capacity. A simple example of such a Markovian model is considered in Section VIII.

Theorem 1: The capacity of the network is:

$$\mu = \frac{p+q}{2d} \quad (1)$$

where

$$p = 1 - \left(1 - \frac{1}{C}\right)^N - \frac{N}{C} \left(1 - \frac{1}{C}\right)^{N-1} \quad (2)$$

$$q = 1 - \left(1 - \frac{1}{C^2}\right)^{N/2} \quad (3)$$

and hence the network can stably support users simultaneously communicating at any rate $\lambda < \mu$.

Note that p represents the probability of finding at least two users in a particular cell, and q represents the probability of finding a source-destination pair within a cell. The proof of the above theorem involves proving that $\lambda \leq \mu$ is necessary for network stability, and that $\lambda < \mu$ is sufficient. Sufficiency is established in Subsection II-C, where a stabilizing algorithm is provided and exact expressions for average delay are derived. A formal proof of necessity is given in Appendix A. Here we provide an abbreviated argument to gain intuition:

Intuitive Explanation: Suppose all users send at rate λ , so that $N\lambda$ represents the sum rate of new packets entering the network. Each of these packets must be transmitted over the network at least once, and are transmitted two or more times if they reach their destinations via a relay node. The maximum rate of single-hop transfers between sources and destinations is Cq (the average number of cells containing a source-destination pair on a given timeslot). All other transmission opportunities must serve packets that take two or more hops to the destination, and the rate of such transmissions is at most $Cp - Cq$ (where Cp is the average number of cells that




Fig. 2. A plot of the limiting capacity $(1 - e^{-d} - de^{-d})/(2d)$ as a function of the user/cell density d .

can support a packet transfer during a given timeslot). Hence: $N\lambda \leq Cq + \frac{Cp-Cq}{2}$, yielding the necessary condition. \square

Taking limits as $N \rightarrow \infty$, we find the network capacity tends to the fixed value $(1 - e^{-d} - de^{-d})/(2d)$. This value tends to zero as d tends either to zero or infinity. Indeed, if d is too large, there will be many users in each cell, most of which will be idle as a single transmitter and receiver are selected. However, if d is too small, the probability of at least two users being in a given cell vanishes. Hence, for nonzero capacity, the ratio $d = N/C$ should be fixed as both N and C scale up. The optimal user/cell density d^* and the corresponding capacity μ^* are: $d^* = 1.7933$, $\mu^* = 0.1492$ (see Fig. 2). Thus, when the number of users N is large, the maximum throughput of a cell partitioned network is close to its limiting maximum of 0.1492 packets/slot. Throughputs arbitrarily close to this value can be achieved by scaling the number of cells C with N to maintain a constant user/cell density d^* .

This μ^* capacity value is close to the maximum throughput estimate of 0.14 packets/slot for the $O(1)$ throughput strategy given by Grossglauser and Tse in [6], where the 0.14 number is obtained by a numerical optimization over a transmit probability θ . In the Grossglauser-Tse strategy, transmitting users send to their nearest neighbors to obtain a high signal to interference ratio on each transmission. The proximity of their optimal throughput to the value of μ^* suggests that when the transmit probability is optimized, the nearest-neighbor transmission policy behaves similarly to a cell-partitioned network. The same value μ^* arises when users send independent data to a finite collection of other users according to a *rate matrix* (λ_{ij}) . In this case, μ^* represents the maximum sum rate into or out of any user provided that no user sends or receives more than any other, as described in Section II-B.

A. Feedback Does not Increase Capacity

We note that the optimal throughput μ of Theorem 1 cannot be improved even if all users have perfect knowledge of future events (see proof of Theorem 1). Thus, control strategies which utilize redundant packet transfers, enable multiple users to overhear the same transmission, or allow for perfect feedback to all users when a given packet has been successfully received, cannot increase capacity.

Corollary 1: The use of redundant packet transfers, multi-user reception, or perfect feedback cannot increase network capacity.

Proof: The capacity region given in Theorem 1 considers all possible strategies, including those which have perfect knowledge of future events. Hence, with full knowledge of the future, any strategy employing redundant packet transfers, multi-user reception, or perfect feedback can be transformed into a policy which does not use these features simply by removing the feedback mechanism (all feedback information would be a-priori known) and deleting all redundant versions of packets, so that only packets which first reach their destination are transmitted. Thus, such features cannot expand the region of stabilizable rates.

However, the capacity region can be *achieved* without feedback, redundancy, or perfect knowledge of the future (as described in the next section) and hence these features do not impact capacity. \square

B. Heterogeneous Demands

Before proceeding with our delay results, we consider the case of communication with heterogeneous rates (λ_{ij}) , where λ_{ij} represents the rate user i receives exogenous data intended for user j . Define the *symmetric capacity region* as the region of all stabilizable data rates such that no user is transmitting or receiving at a higher total data rate than any other. Let K represent the maximum number of destination users to which a source transmits (i.e., for each user i , at most K of the λ_{ij} terms are nonzero).

Theorem 2: The symmetric capacity region of the network has the form:

$$\sum_j \lambda_{ij} \leq \frac{(1 - e^{-d} - de^{-d})}{2d} + O(K/N) \quad \forall i \quad (4)$$

$$\sum_j \lambda_{ij} \leq \frac{(1 - e^{-d} - de^{-d})}{2d} + O(K/N) \quad \forall j \quad (5)$$

Proof: This proof is similar to the proof of Theorem 1, and the differences are given in Appendix C. \square

In the next section we present a capacity achieving strategy together with an exact delay analysis. To simplify the discussion, throughout the rest of this paper we assume that each user communicates with rate λ to a unique destination according to the pairing $1 \leftrightarrow 2, 3 \leftrightarrow 4$, etc., so that $K = 1$ and the exact capacity result $\mu = (p + q)/(2d)$ of Theorem 1 applies for all network sizes N .

C. Delay Analysis and the 2-Hop Relay Algorithm

In this section, we consider a modified version of the Grossglauser-Tse relay algorithm of [6], and show the

algorithm is capacity achieving with a bounded average delay. The algorithm restricts packets to 2-hop paths, where on the first hop a packet is transmitted to any available user. This user will act as a ‘relay’ for the packet. The packet is stored in the buffer of the relay until an opportunity arises for it to be transmitted by the relay to its destination. Note that the notion of relaying is vitally important, as it allows throughput to be limited only by the rate at which a source encounters other users, rather than by the rate at which a source encounters its destination.

Cell Partitioned Relay Algorithm: Every timeslot and for each cell containing at least two users:

- 1) If there exists a source-destination pair within the cell, randomly choose such a pair (uniformly over all such pairs in the cell). If the source contains a new packet intended for that destination, transmit. Else remain idle.
- 2) If there is no source-destination pair in the cell, designate a random user within the cell as sender. Independently choose another user as receiver among the remaining users within the cell. With equal probability, randomly choose one of the two options:
 - *Send a Relay packet to its Destination:* If the designated transmitter has a packet destined for the designated receiver, send that packet to the receiver. Else remain idle.
 - *Send a New Relay Packet:* If the designated transmitter has a new packet (one that has never before been transmitted), relay that packet to the designated receiver. Else remain idle.

Since packets that have already been relayed are restricted from being transmitted to any user other than their destination, the above algorithm restricts all routes to 2-hop paths. The algorithm schedules packet transfer opportunities without considering queue backlog. Performance can be improved by allowing alternative scheduling opportunities in the case when no packet is available for the chosen transmission. However, the randomized nature of the algorithm admits a nice *decoupling* between sessions (see Fig. 3), where individual users see the network only as a source, destination, and intermediate relays, and transmissions of packets for other sources are reflected simply as random ON/OFF service opportunities.

Theorem 3: Consider a cell partitioned network (with N users and C cells) under the 2-hop relay algorithm, and assume that users change cells *i.i.d.* and uniformly over each cell every timeslot. If the exogenous input stream to user i is a Bernoulli stream of rate λ_i (where $\lambda_i < \mu$), then the total network delay W_i for user i traffic satisfies:

$$\mathbb{E}\{W_i\} = \frac{N-1-\lambda_i}{\mu-\lambda_i} \quad (6)$$

where the capacity μ is defined in (1).

Proof: The proof uses reversibility of the first stage queue, and is provided in Appendix B. \square

Note that the decoupling property of the cell partitioned relay algorithm admits a decoupled delay bound, so that the waiting time for user i packets depends only on the rate of the input stream for user i , and does not depend on the

rate of other streams—even if the rate of these streams is greater than capacity. It follows that the network is stable with bounded delays whenever all input streams are less than capacity, i.e., when $\lambda_i < \mu$ for all users i . Thus, the relay algorithm achieves the capacity bound given in (1) of Theorem 1. It is perhaps counter-intuitive that the algorithm achieves capacity, as it often forces cells to remain idle even when choosing an alternate sender would allow for a packet to be delivered to its destination. The intuition is that all cases of idleness arise because a queue is empty, an event that becomes increasingly unlikely as load approaches capacity.

The form of the delay expression is worth noting. First note the classic $1/(\mu - \lambda_i)$ behavior, representing the asymptotic growth in delay as data rates are pushed towards the capacity boundary. Second, note that for a fixed loading value $\rho_i = \lambda_i/\mu$, delay is $O(N)$, growing linearly in the size of the network.

The exact delay analysis is enabled by the Bernoulli input assumption. If inputs are assumed to be Poisson, the delay theory of [1] [5] can be used to develop a delay bound, and the bound for Poisson inputs is not considerably different from the exact expression for Bernoulli inputs given in (6). These results can also be extended to the case when the mobility model conforms to a Markovian random walk (see analytical discussion and simulation results in Sections VII and VIII).

III. SENDING A SINGLE PACKET

In the previous subsection we showed that the cell partitioned relay algorithm yields an average delay of $O(N/(\mu - \lambda_i))$. Inspection of (6) shows that this $O(N)$ characteristic cannot be removed by decreasing the data rate λ . The following questions emerge: Can another scheduling algorithm be constructed which improves delay? What is the minimum delay the network can guarantee, and for what data rates is this delay obtainable? More generally, for a given data rate λ (assumed to be less than the system capacity μ), we ask: What is the optimal delay bound, and what algorithm achieves this? In this section we present several fundamental bounds on delay performance, which establishes initial steps towards addressing these general questions. We assume throughout that

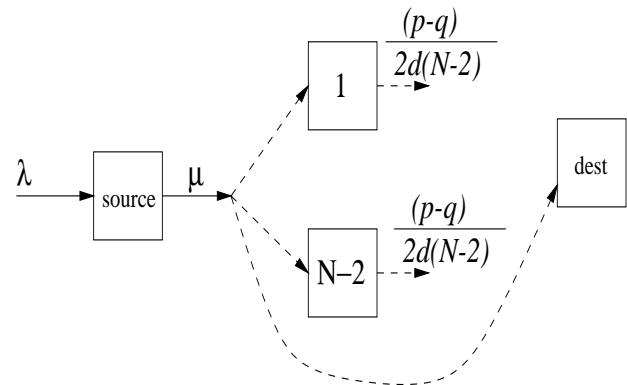


Fig. 3. A decoupled diagram of the network as seen by the packets transmitted from a single user to the corresponding destination. Service opportunities at the first stage are Bernoulli with rate μ . Service at the second stage (relay) queues is Bernoulli with rate $(p - q)/(2d(N - 2))$.

the user/cell density d is a fixed value independent of N , and use $d = d^* = 1.7933$ in all numerical examples.

A. Scheduling Without Redundancy

Suppose that no redundancy is used: that is, packets are not duplicated and are held by at most one user of the network at any given time. Thus, a packet that is transmitted to another user is deleted from the memory storage of the transmitting user. Note that this is the traditional approach to data networking, and that the 2-hop relay algorithm is in this class of algorithms.

Theorem 4: Algorithms which do not use redundancy cannot achieve an average delay of less than $O(N)$.

Proof: The minimum delay of any packet is computed by considering the situation where the network is empty and user 1 sends a single packet to user 2. It is easy to verify that relaying the packet cannot help, and hence the delay distribution is geometric with mean $C = N/d$. \square

Hence, the relay algorithm not only achieves capacity, but achieves the optimal $O(N)$ delay performance among all strategies which do not use redundancy. Other policies which do not use redundancy can perhaps improve upon the delay coefficient, but cannot change the $O(N)$ characteristic.

B. Scheduling With Redundancy

Although redundancy cannot increase capacity, it can considerably improve delay. Clearly, the time required for a packet to reach the destination can be reduced by repeatedly transmitting this packet to many users of the network—improving the chances that some user holding an original or duplicate version of the packet reaches the destination. Consider any network algorithm (which may or may not use redundant packet transfers) that restricts packets to 2-hop paths.

Theorem 5: No algorithm (with or without redundancy) which restricts packets to 2-hop paths can provide an average delay better than $O(\sqrt{N})$.

To prove this result, again consider the sending of a single packet from its source to its destination. Clearly the optimal scheme is to have the source send duplicate versions of the packet to new relays whenever possible, and for the packet to be relayed to the destination as soon as either the source or a duplicate-carrying relay enters the same cell as the destination. Let T_N represent the time required to reach the destination under this optimal policy for sending a single packet. In the following lemma, we bound the limiting behavior¹ of $\mathbb{E}\{T_N\}$, proving Theorem 5.

Lemma 1: $e^{-d} \leq \lim_{N \rightarrow \infty} \frac{\mathbb{E}\{T_N\}}{\sqrt{N}} \leq \frac{2}{1-e^{-d}}$

Proof: Lemma 1 (a) *Lower Bound:* To prove the lower bound, note that during timeslots $\{1, 2, \dots, \sqrt{N}\}$, there are fewer than \sqrt{N} users holding the packet. Hence, $Pr[T_N > \sqrt{N}] \geq (1 - 1/C)^{\sqrt{N}\sqrt{N}}$ (where $(1 - 1/C)^{\sqrt{N}}$ is the probability that nobody within a group of \sqrt{N} particular users

enters the cell of the destination during a given timeslot). Recall that the user/cell density d is defined $d \triangleq N/C$. Thus:

$$\begin{aligned} \mathbb{E}\{T_N\} &\geq \mathbb{E}\{T_N | T_N > \sqrt{N}\} Pr[T_N > \sqrt{N}] \\ &\geq \sqrt{N} \left(1 - \frac{d}{N}\right)^N \rightarrow e^{-d} \sqrt{N} \end{aligned}$$

(b) *Upper Bound:* To prove the upper bound, note that $\mathbb{E}\{T_N\} \leq S_1 + S_2$, where S_1 represents the expected number of slots required to send out duplicates of the packet to \sqrt{N} different users, and S_2 represents the expected time until one user within a group of \sqrt{N} users containing the packet reaches the cell of the destination. The probability of the source meeting a new user is at least $1 - (1 - 1/C)^{N-\sqrt{N}}$ for every timeslot where fewer than \sqrt{N} users have packets, and hence the average time to reach a new user is less than or equal to the inverse of this quantity (i.e., the average time of a geometric variable). As the source must encounter \sqrt{N} users, we have:

$$S_1 \leq \frac{\sqrt{N}}{1 - (1 - 1/C)^{N-\sqrt{N}}} \rightarrow \frac{\sqrt{N}}{1 - e^{-d}}$$

To compute S_2 , note that $P(\text{success})$, the probability that one of the \sqrt{N} users reaches the destination during a slot, is given by the probability there is at least one other user in the same cell as the destination multiplied by the conditional probability that a packet-carrying user is present given there is at least one other user in the cell. The former probability is $1 - (1 - 1/C)^{N-1}$, and the latter is at least \sqrt{N}/N :

$$P(\text{success}) \geq \frac{1 - (1 - 1/C)^{N-1}}{\sqrt{N}} \rightarrow \frac{1 - e^{-d}}{\sqrt{N}} \quad (7)$$

Hence, $S_2 \leq \frac{\sqrt{N}}{1 - e^{-d}}$. Summing S_1 and S_2 proves the result. \square

An exact expression for the minimum delay $\mathbb{E}\{T_N\}$ is presented in Appendix D by using a recursive formula. In Fig. 4 we plot the exact expression as a function of N together with the upper and lower bounds of Lemma 1 for the case $d = d^* = 1.7933$.

C. Multi-User Reception

To increase the packet replication speed throughout the network, it is useful to allow a transmitted packet to be received by *all other users* in the same cell as the transmitter, not just the single intended recipient. This feature cannot increase capacity, but can considerably improve delay by enabling multiple duplicates to be injected into the network with just a single transmission. However, the $O(\sqrt{N})$ result of Theorem 5 cannot be overcome by introducing multi-user reception (see Appendix E). For the remainder of this paper, we assume multi-user reception is available when proving fundamental performance limits, but we do not require multi-user reception in any of our algorithms that demonstrate achievability of these limits.

¹Using the inequality $e^{-\frac{d^2}{N-d}} e^{-d} \leq \left(1 - \frac{d}{N}\right)^N \leq e^{-d}$, explicit bounds of the form $\alpha\sqrt{N} \leq \mathbb{E}\{T_N\} \leq \beta\sqrt{N}$ can also be derived.

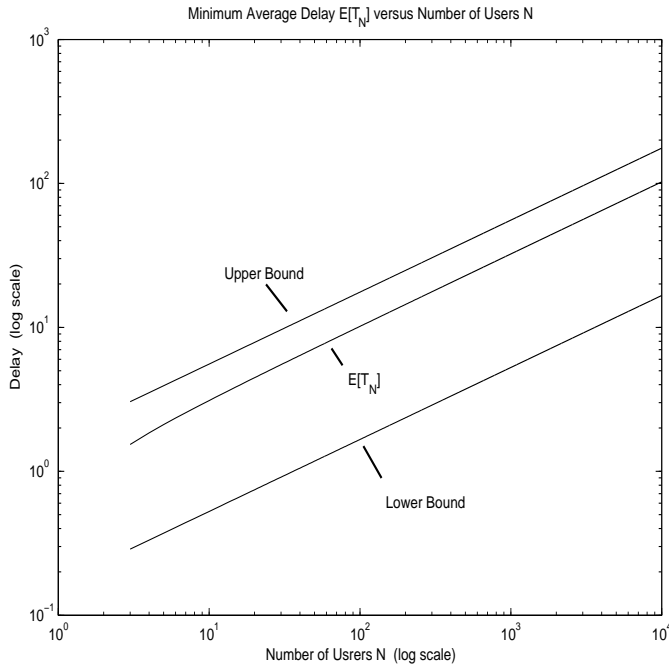


Fig. 4. The exact minimum delay of a 2-hop scheduling scheme versus the number of users N at the optimal user/cell density d^* , together with the upper and lower bounds of Lemma 1. Curves are plotted on a log-log scale and have slope $1/2$, illustrating the $O(\sqrt{N})$ behavior.

IV. SCHEDULING FOR DELAY IMPROVEMENT

In the previous section an $O(\sqrt{N})$ delay bound was developed for redundant scheduling by considering a single packet for a single destination. Two complications arise when designing a general scheduling protocol using redundancy: (1) All sessions must use the network simultaneously, and (2) Remnant versions of a packet that has already been delivered to its destination create excess congestion and must somehow be removed.

Here we show that the properties of the 2-hop relay algorithm make it naturally suited to treat the multi-user problem. The second complication of excess packets is overcome by the following *in-cell feedback protocol*, in which a receiving node tells its transmitter which packet it is looking for before transmission begins. We assume all packets are labeled with *send numbers* SN , and the in-cell feedback is in the form of a *request number* RN delivered by the destination to the transmitter just before transmission. In the following protocol, each packet is retransmitted \sqrt{N} times to distinct relay users.

In-Cell Feedback Scheme with \sqrt{N} Redundancy: In every cell with at least two users, a random sender and a random receiver are selected, with uniform probability over all users in the cell. With probability $1/2$, the sender is scheduled to operate in either ‘source-to-relay’ mode, or ‘relay-to-destination’ mode, described as follows:

- 1) *Source-to-Relay Mode:* The sender transmits packet SN , and does so upon every transmission opportunity until \sqrt{N} replicas have been delivered to distinct users, or until the sender transmits SN directly to the destination. After

such a time, the send number is incremented to $SN + 1$. If the sender does not have a new packet to send, remain idle.

- 2) *Relay-to-Destination Mode:* When a user is scheduled to transmit a relay packet to its destination, the following handshake is performed:

- The receiver delivers its current RN number for the packet it desires.
- The transmitter deletes all packets in its buffer destined for this receiver which have SN numbers lower than RN .
- The transmitter sends packet RN to the receiver. If the transmitter does not have the requested packet RN , it remains idle for that slot.

Notice that the destination receives all packets in order, and that no packet is ever transmitted twice to its destination.

Theorem 6: The In-Cell Feedback Scheme achieves the $O(\sqrt{N})$ delay bound, with user data rates of $O(1/\sqrt{N})$.

More precisely, if all users receive exogenous data for their destinations according to a Poisson process of rate λ_i , the network can stably support rates $\lambda_i < \tilde{\mu}$, for the reduced network throughput $\tilde{\mu}$ given by:

$$\tilde{\mu} = \frac{\gamma_N (1 - e^{-d})}{4(2 + d)\sqrt{N}} \quad (8)$$

where γ_N is a sequence that converges to 1 as $N \rightarrow \infty$. Furthermore, average end-to-end delay $\mathbb{E}\{W_i\}$ satisfies:

$$\mathbb{E}\{W_i\} \leq \frac{1}{2} + \frac{1/\tilde{\mu}}{1 - \rho_i}$$

where $\rho_i \triangleq \lambda_i / \tilde{\mu}$.

To prove the result, first note that when a new packet reaches the head of the line at its source queue, the time required for the packet to reach its destination is at most $T_N = S_1 + S_2$, where S_1 represents the time required for the source to send out \sqrt{N} replicas of the packet, and S_2 represents the time required to reach the destination given that \sqrt{N} users have the packet. Bounds on the expectations of S_1 and S_2 which are independent of the initial state of the network can be computed in a manner similar to the proof of Lemma 1. The multi-user environment here simply acts to scale up these expectations by a constant factor due to collisions with other users (compare the upper bound of Lemma 1 with that given in (9) below). This factor does not scale with N because the average number of users in any cell is the finite number d . Indeed, in Appendix F it is shown that:

$$\mathbb{E}\{T_N\} \leq \frac{4(2 + d)\sqrt{N}}{\gamma_N(1 - e^{-d})} \quad (9)$$

where γ_N is a function that converges to 1 as $N \rightarrow \infty$.

Note that the random variable T_N satisfies the *sub-memoryless* property: The residual time of T_N given that a fixed number of slots have already passed (without T_N expiring) is stochastically less than the original time T_N .² This is because the topology of the network is independent from slot to slot, and hence starting out with several duplicate packets

²This is often called the ‘New Better than Used’ property’, see [19].

already in the network yields statistically smaller delay than if no such initial duplicates are present.

The RN/SN handshake ensures that newer packets do not interfere with older packets, but that replication of the next packet waiting at the source queue begins on or before completion of the T_N ‘service time’ for the current packet SN . Packets thus view the network as a single queue to which they arrive and are served sequentially. Although actual service times may not be i.i.d., their expectations are all independently bounded by $\mathbb{E}\{T_N\}$, as are the expected residual service times seen by a randomly arriving packet. This is sufficient to establish the following lemma, the proof of which is similar to the derivation of the standard P-K formula for average delay in an $M/G/1$ queue.

Lemma 2: Suppose inputs to a single server queue are Poisson with sub-memoryless service times that are independently bounded in expectation by a value $\mathbb{E}\{T_N\}$. If the arrival rate is λ , where $\lambda < 1/\mathbb{E}\{T_N\}$, then average delay satisfies:

$$\mathbb{E}\{W\} \leq \frac{1}{2} + \frac{\mathbb{E}\{T_N\}}{1 - \rho} \quad (10)$$

where $\rho \triangleq \lambda \mathbb{E}\{T_N\}$. The expression on the right hand side of the above inequality is the standard expression for delay in an $M/M/1$ queue with i.i.d. service times T_N that are restricted to start on slot boundaries.

Proof: Consider a single packet arriving from a Poisson stream, and let W_q represent the time this packet spends waiting in the queue before reaching the server. We have:

$$W_q = \sum_{i=1}^{N_q} X_i + R \quad (11)$$

where N_q is the number of packets already in the queue, $\{X_i\}$ are the service times of these packets, and R represents the residual time until either the packet currently in the server finishes its service, or (if the system is empty) the start of a new timeslot. Note that $\mathbb{E}\{R\} \leq \rho_{actual} \mathbb{E}\{T_N\} + (1 - \rho_{actual})\frac{1}{2}$, where ρ_{actual} represents the probability that the system is busy with a packet already in service. From Little’s Theorem we have that $\rho_{actual} = \lambda \mathbb{E}\{X\}$, where $\mathbb{E}\{X\}$ represents the average service time of a generic packet. Since $\mathbb{E}\{X\} \leq \mathbb{E}\{T_N\}$, it follows that $\rho_{actual} \leq \rho$. Clearly $\mathbb{E}\{T_N\} \geq 1/2$, and hence we can increase the upper bound on $\mathbb{E}\{R\}$ by replacing ρ_{actual} with ρ , yielding $\mathbb{E}\{R\} \leq \rho \mathbb{E}\{T_N\} + (1 - \rho)\frac{1}{2}$. Taking expectations of (11) thus yields:

$$\begin{aligned} \mathbb{E}\{W_q\} &= \mathbb{E}_{N_q} \left[\sum_{i=1}^{N_q} \mathbb{E}\{X_i | N_q\} \right] + \mathbb{E}\{R\} \\ &\leq \mathbb{E}_{N_q} \left[\sum_{i=1}^{N_q} \mathbb{E}\{T_N\} \right] + \rho \mathbb{E}\{T_N\} + (1 - \rho)\frac{1}{2} \\ &= \mathbb{E}\{N_q\} \mathbb{E}\{T_N\} + \rho \mathbb{E}\{T_N\} + (1 - \rho)\frac{1}{2} \\ &= \lambda \mathbb{E}\{W_q\} \mathbb{E}\{T_N\} + \rho \mathbb{E}\{T_N\} + (1 - \rho)\frac{1}{2} \end{aligned} \quad (12)$$

where (12) follows from Little’s Theorem. We thus have:

$$\mathbb{E}\{W_q\} \leq \frac{\rho \mathbb{E}\{T_N\}}{1 - \rho} + \frac{1}{2}$$

Noting that the total waiting time $\mathbb{E}\{W\}$ satisfies $\mathbb{E}\{W\} \leq \mathbb{E}\{W_q\} + \mathbb{E}\{T_N\}$ yields the result. \square

Defining $\tilde{\mu} \triangleq 1/\mathbb{E}\{T_N\}$ and using (9) proves Theorem 6.

V. MULTI-HOP SCHEDULING AND LOGARITHMIC DELAY

To further improve delay, we can remove the 2-hop restriction and consider schemes which allow for multi-hop paths. Here, a simple flooding protocol is developed and shown to achieve $O(\log(N))$ delay at the expense of further reducing throughput.

To achieve $O(\log(N))$ delay, consider the situation in which a single packet is delivered over an empty network. At first, only the source user contains the packet. The packet is transmitted and received by all other users in the same cell as the source. In the next timeslot, the source as well as all of the new users containing the packet transmit in their respective cells, and so on. If all duplicate-carrying users enter distinct cells every timeslot, and each of these users delivers the packet to exactly one new user, then the number of users containing the packet grows geometrically according to the sequence $\{1, 2, 4, 8, 16, \dots\}$. The actual growth pattern may deviate from this geometric sequence somewhat, due to multiple users entering the same cell, or to users entering cells that are devoid of other users. However, it can be shown that the *expected* growth is geometric provided that the number of packet-holding users is less than $N/2$.

Define the total time to reach all users as $T_N = S_1 + S_2$, where S_1 and S_2 respectively represent the time required to send the packet to at least $N/2$ users, and the time required to deliver the packet to the remaining users given that at least $N/2$ users initially hold the packet.

Lemma 3: Under the above algorithm of flooding the network with a single packet, for any network size $N \geq \max\{d, 2\}$, the expected time $\mathbb{E}\{T_N\}$ for the packet to reach every user satisfies $\mathbb{E}\{T_N\} \leq \mathbb{E}\{S_1\} + \mathbb{E}\{S_2\}$, where:

$$\begin{aligned} \mathbb{E}\{S_1\} &\leq \frac{\log(N)(1 + d/2)}{\log(2)(1 - e^{-d/2})} \\ \mathbb{E}\{S_2\} &\leq 1 + \frac{2}{d}(1 + \log(N/2)) \end{aligned} \quad (13)$$

Proof: The proof is given in Appendix G. \square

A. Fair Packet Flooding Protocol

Thus, $O(\log(N))$ delay is achievable when sending a single packet over an empty network. To enable $O(\log(N))$ delay in the general case where all sessions are active and share the network resources, we construct a flooding protocol in which the oldest packet that has not been delivered to all users is selected to dominate network resources. We assume that packets are sequenced with SN numbers as before. Additionally, packets are stamped with the timeslot t in which they arrived.

Fair Packet Flooding Protocol: Every timeslot and in each cell, users perform the following: Among all packets contained in at least one user of the cell but which have never been received by some other user in the same cell, choose the packet p which arrived earliest (i.e., it has the smallest timestamp t_p). If there are ties, choose the packet from the session i which

maximizes $(t_p + i) \bmod N$. Transmit this packet to all other users in the cell. If no such packet exists, remain idle.

The above protocol is ‘fair’ in that in case of ties, session i packets are given top priority every N timeslots. Other schemes for choosing which packet to dominate the network could also be considered. Delay under the above protocol can be understood by comparing the network to a single queue with N input streams of rates $\lambda_1, \lambda_2, \dots, \lambda_N$ which share a single server with service times T_N . Note that the T_N service time is also sub-memoryless. Thus, from Lemma 2, we have:

Theorem 7: For Poisson inputs with rates λ_i for each source i , the network under the fair flooding protocol is stable whenever $\sum_i \lambda_i < 1/\mathbb{E}\{T_N\}$, with average end-to-end delay satisfying:

$$\mathbb{E}\{W\} \leq \frac{1}{2} + \frac{\mathbb{E}\{T_N\}}{1 - \rho} \quad (14)$$

where $\rho \triangleq \sum_i \lambda_i \mathbb{E}\{T_N\}$, and $\mathbb{E}\{T_N\} = \mathbb{E}\{S_1\} + \mathbb{E}\{S_2\}$. Note that $O(\log(N))$ bounds on $\mathbb{E}\{S_1\}$ and $\mathbb{E}\{S_2\}$ are given in Lemma 3. Thus, when all sources have identical input rates λ , stability and logarithmic delay is achieved when $\lambda = O(\frac{1}{N \log(N)})$. \square

Note that the flooding algorithm easily allows for *multicast sessions*, where data of rate λ is delivered from each source to *all other users*. One might expect that delay can be improved if we only design for unicast. However, it is shown in Appendix H that logarithmic delay is the best possible for any strategy at any data rate. Hence, communication for unicast or multicast is the same in the logarithmic delay regime. In the next section, we address the following question: Is it possible to increase data rates via some other protocol while maintaining the same average delay guarantees?

VI. FUNDAMENTAL DELAY/RATE TRADEOFFS

Considering the capacity achieving 2-hop relay algorithm, the 2-hop algorithm with \sqrt{N} redundancy, and the packet flooding protocol, we have the following achievable delay/capacity performance tradeoffs.

scheme	throughput	delay
no redundancy	$O(1)$	$O(N)$
redundancy 2-hop	$O(1/\sqrt{N})$	$O(\sqrt{N})$
redundancy multi-hop	$O(\frac{1}{N \log(N)})$	$O(\log(N))$

A simple observation reveals that $\text{delay}/\text{rate} \geq O(N)$ for each of these three protocols. In this section, we establish that this is in fact a necessary condition. Thus, performance of each given protocol falls on or near the boundary of a *fundamental rate-delay curve* (see Fig. 1).

Consider a network with N users, and suppose all users receive packets at the same rate λ . A control protocol which makes decisions about scheduling, routing, and packet retransmissions is used to stabilize the network and deliver all packets to their destinations while maintaining an average end-to-end delay less than some threshold \bar{W} .

Theorem 8: A necessary condition for any conceivable routing and scheduling protocol that stabilizes the network

with input rates λ while maintaining bounded average end-to-end delay \bar{W} is given by:

$$\frac{\bar{W}}{\lambda} \geq \frac{N - d}{4d} (1 - \log(2)) \quad (15)$$

where $\log()$ denotes the natural logarithm, and $d = N/C$ is the user/cell density.

In particular, if $d = O(1)$, then $\bar{W}/\lambda \geq O(N)$.

The above condition holds for all possible control strategies, including those that use multi-user reception. We prove this theorem with a novel technique for probabilistic conditioning.

Proof: Suppose the input rate of each of the N sessions is λ , and there exists some stabilizing scheduling strategy which ensures an end-to-end delay of \bar{W} . In general, the end-to-end delay of packets from individual sessions could be different, and we define \bar{W}_i as the resulting average delay of packets from session i . We thus have:

$$\bar{W} = \frac{1}{N} \sum_i \bar{W}_i \quad (16)$$

Let \bar{R}_i represent the average *redundancy* associated with packets from session i . That is, \bar{R}_i is the number of users who receive a copy of an individual packet during the course of the network control operation, averaged over all packets from session i . Note that all packets are eventually received by the destination, so that $\bar{R}_i \geq 1$. Additional redundancy could be introduced by multi-hop routing, or by any packet replication effort that is used to achieve stability and/or improve delay. The average number of successful packet receptions per timeslot is thus given by the quantity $\lambda \sum_{i=1}^N \bar{R}_i$. Since each of the N users can receive at most 1 packet per timeslot, we have:

$$\lambda \sum_{i=1}^N \bar{R}_i \leq N \quad (17)$$

Now consider a single packet p which enters the network from session i . This packet has an average delay of \bar{W}_i and an average redundancy of \bar{R}_i . Let random variables W_i and R_i represent the actual delay and redundancy for this packet. We have:

$$\bar{W}_i = \mathbb{E}\{W_i \mid R_i \leq 2\bar{R}_i\} Pr[R_i \leq 2\bar{R}_i] + \quad (18)$$

$$\mathbb{E}\{W_i \mid R_i > 2\bar{R}_i\} Pr[R_i > 2\bar{R}_i]$$

$$\geq \mathbb{E}\{W_i \mid R_i \leq 2\bar{R}_i\} Pr[R_i \leq 2\bar{R}_i]$$

$$\geq \mathbb{E}\{W_i \mid R_i \leq 2\bar{R}_i\} \frac{1}{2} \quad (19)$$

where (19) follows because $Pr[R_i \leq 2\bar{R}_i] \geq \frac{1}{2}$ for any non-negative random variable R_i .

Note that the smallest possible delay for packet p is the time required for one of its carriers to enter the same cell as the destination. Consider now a virtual system in which there are $2\bar{R}_i$ users initially holding packet p , and let Z represent the time required for one of these users to enter the same cell as the destination. Every timeslot the ‘success probability’ for this system is $\phi \triangleq 1 - (1 - \frac{1}{C})^{2\bar{R}_i}$, so that $\mathbb{E}\{Z\} = 1/\phi$. Although there are more users holding packet p in this system, the expectation of Z does not necessarily bound $\mathbb{E}\{W_i \mid R_i \leq 2\bar{R}_i\}$ because conditioning on the event

$\{R_i \leq 2\bar{R}_i\}$ might skew the probabilities associated with the user mobility process. However, since the event $\{R_i \leq 2\bar{R}_i\}$ occurs with probability at least $1/2$, we obtain the following bound:

Claim 1:

$$\mathbb{E}\{W_i \mid R_i \leq 2\bar{R}_i\} \geq \inf_{\Theta} \mathbb{E}\{Z \mid \Theta\} \quad (20)$$

where the conditional expectation is minimized over all conceivable events Θ which occur with probability greater than or equal to $1/2$.

A proof of Claim 1 is given at the end of this subsection. Intuitively, it holds because minimizing over all such events Θ includes events that could yield mobility patterns of the type encountered when $\{R_i \leq 2\bar{R}_i\}$.

We now stochastically couple Z to an independent exponential variable \tilde{Z} with rate $\gamma \triangleq \log(1/(1-\phi))$. The variable \tilde{Z} is stochastically less than Z because $Pr[\tilde{Z} > \omega] \leq Pr[Z > \omega]$ for all ω . Indeed, because \tilde{Z} is exponential with rate γ , we have $Pr[\tilde{Z} > \omega] = e^{-\gamma\omega} = (1-\phi)^\omega$ for any $\omega \geq 0$, while Z is geometric with success probability ϕ , so that:

$$Pr[Z > \omega] = (1-\phi)^{\lfloor \omega \rfloor} \geq (1-\phi)^\omega = Pr[\tilde{Z} > \omega]$$

The fact that \tilde{Z} is stochastically less than Z leads to the following claim:

Claim 2: For variables Z and \tilde{Z} , we have:

$$\inf_{\Theta} \mathbb{E}\{Z \mid \Theta\} \geq \inf_{\tilde{\Theta}} \mathbb{E}\{\tilde{Z} \mid \tilde{\Theta}\} = \frac{1 - \log(2)}{\gamma} \quad (21)$$

where the first infimum is taken over all events Θ that occur with probability greater than or equal to $1/2$ on the probability space for Z , and the second infimum is taken over all events $\tilde{\Theta}$ that occur with probability greater than or equal to $1/2$ on the probability space for \tilde{Z} .

Claim 2 is proven at the end of this subsection. Using (21) and (20) in (19) yields: $\bar{W}_i \geq \frac{1 - \log(2)}{2\gamma}$. From the definitions of γ and ϕ , we have $\gamma = \log\left(1/(1 - \frac{1}{C})^{2\bar{R}_i}\right) = 2\bar{R}_i \log(1 + \frac{1}{C-1})$. Since $\log(1+x) \leq x$ for any x , we have $\gamma \leq 2\bar{R}_i/(C-1)$. We thus have:

$$\bar{W}_i \geq \frac{1 - \log(2)}{2\gamma} \geq \frac{(C-1)(1 - \log(2))}{4\bar{R}_i}$$

Summing this inequality over all i , we have:

$$\begin{aligned} \bar{W} &= \frac{1}{N} \sum_{i=1}^N \bar{W}_i \geq \frac{(C-1)(1 - \log(2))}{4} \frac{1}{N} \sum_{i=1}^N \frac{1}{\bar{R}_i} \\ &\geq \frac{(C-1)(1 - \log(2))}{4 \frac{1}{N} \sum_{i=1}^N \bar{R}_i} \end{aligned} \quad (22)$$

where (22) follows from Jensen's inequality, noting that the function $f(R) = 1/R$ is convex, and hence $\frac{1}{N} \sum_{i=1}^N f(\bar{R}_i) \geq f(\frac{1}{N} \sum_{i=1}^N \bar{R}_i)$. Combining (22) and (17), we have:

$$\bar{W} \geq \frac{(C-1)(1 - \log(2))\lambda}{4} = \frac{(N-d)(1 - \log(2))\lambda}{4d}$$

Hence, the delay/rate characteristics necessarily satisfy the inequality $\frac{\bar{W}}{\lambda} \geq O(N)$, proving the theorem. \square

We complete the analysis by proving Claims 1 and 2:

Proof of Claim 2. We first compute $\inf_{\tilde{\Theta}} \mathbb{E}\{\tilde{Z} \mid \tilde{\Theta}\}$. Note that \tilde{Z} is a continuous variable, and so the minimizing event $\tilde{\Theta}$ is clearly the event $\{\tilde{Z} \leq \omega\}$, where ω is the smallest value such that $Pr[\tilde{Z} \leq \omega] \geq \frac{1}{2}$. Since \tilde{Z} is exponential with rate $\gamma = \log(1/(1-\phi))$, we have $Pr[\tilde{Z} > \omega] = e^{-\gamma\omega} = 1/2$, and hence $\omega = \frac{\log(2)}{\gamma}$. Conditioning on this event, we have:

$$\begin{aligned} \inf_{\tilde{\Theta}} \mathbb{E}\{\tilde{Z} \mid \tilde{\Theta}\} &= \mathbb{E}\{\tilde{Z} \mid \tilde{Z} \leq \omega\} \\ &= \frac{\mathbb{E}\{\tilde{Z}\} - \mathbb{E}\{\tilde{Z} \mid \tilde{Z} > \omega\} Pr[\tilde{Z} > \omega]}{Pr[\tilde{Z} \leq \omega]} \\ &= \frac{\frac{1}{\gamma} - (\omega + \frac{1}{\gamma})\frac{1}{2}}{1/2} = \frac{1 - \log(2)}{\gamma} \end{aligned}$$

Now note that \tilde{Z} is stochastically less than Z , so that there must exist a coupling variable Z' such that variables \tilde{Z} and Z' have the same distribution, and Z' lies on the same probability space as Z and satisfies $Z' \leq Z$ for all instances of Z' and Z (see [19] for a discussion of stochastic coupling). Since Z' is also an exponential with rate γ , it follows that $\inf_{\Theta} \mathbb{E}\{Z' \mid \Theta\} = (1 - \log(2))/\gamma$. However, because $Z' \leq Z$ always, it follows that:

$$\inf_{\Theta} \mathbb{E}\{Z' \mid \Theta\} \leq \inf_{\Theta} \mathbb{E}\{Z \mid \Theta\}$$

proving Claim 2. \square

To prove Claim 1, we present a preliminary lemma:

Lemma 4: For any random variables X, Y such that X is stochastically greater than Y , and for any event Φ such that $Pr[\Phi] \geq 1/2$ (where Φ occurs on the same probability space as X), we have:

$$\mathbb{E}\{X \mid \Phi\} \geq \inf_{\{\Theta \mid Pr[\Theta] \geq \frac{1}{2}\}} \mathbb{E}\{Y \mid \Theta\}$$

where events Θ occur on the same probability space as Y .

Proof of Lemma 4. Since X is stochastically greater than Y , there must exist a variable \tilde{X} defined on the same probability space as X , such that $X \geq \tilde{X}$ always, and where \tilde{X} and Y have the same distribution [19]. Thus:

$$\begin{aligned} \mathbb{E}\{X \mid \Phi\} &\geq \mathbb{E}\{\tilde{X} \mid \Phi\} \\ &\geq \inf_{\{\Psi \mid Pr[\Psi] \geq \frac{1}{2}\}} \mathbb{E}\{\tilde{X} \mid \Psi\} \end{aligned} \quad (23)$$

$$= \inf_{\{\Theta \mid Pr[\Theta] \geq \frac{1}{2}\}} \mathbb{E}\{Y \mid \Theta\} \quad (24)$$

where (23) follows because Φ is a particular element of the collection of sets Ψ that occur on the same probability space as \tilde{X} and satisfy $Pr[\Psi] \geq \frac{1}{2}$, and (24) follows because \tilde{X} and Y have the same distribution, and the value of any such infimum depends only on the distribution (see Appendix D). \square

Proof of Claim 1: Recall that W_i represents the delay of packet p under a general scheduling strategy, and R_i represents the redundancy associated with scheduling this packet. Let W_i^{rest} represent the corresponding delay under the restricted scheduling policy that schedules packets as before until either the packet is successfully delivered, or

the redundancy increases to $2\bar{R}_i$ (at which point no more redundant transmissions are allowed). Since this modified policy restricts redundancy to at most $2\bar{R}_i$, the delay W_i^{rest} is *stochastically greater* than the variable Z , representing the delay in a virtual system with only one packet that is initially held by $2\bar{R}_i$ users. By Lemma 4 we thus have that $\mathbb{E}\{W_i^{rest} | R_i \leq 2\bar{R}_i\} \geq \inf_{\{\Theta | Pr[\Theta] \geq \frac{1}{2}\}} \mathbb{E}\{Z | \Theta\}$. However, note that the restricted policy is identical to the original policy whenever $R_i \leq 2\bar{R}_i$, and hence $\mathbb{E}\{W_i | R_i \leq 2\bar{R}_i\} = \mathbb{E}\{W_i^{rest} | R_i \leq 2\bar{R}_i\}$, proving Claim 1. \square

A. Discussion

The fact that $\text{delay}/\text{rate} \geq O(N)$ establishes a fundamental performance tradeoff, illustrating that no scheduling and routing algorithm can simultaneously yield low delay and high throughput. The $O(N)$ and $O(\sqrt{N})$ scheduling algorithms provided here meet this bound with equality, and the $O(\log(N))$ algorithm lies above the bound by a factor of $O(\log^2(N))$ (see table above).

We note that alternate approaches to the capacity/delay tradeoff problem were recently developed in [17] [16] [18] for networks with different physical characteristics. Specifically, the work in [17] develops a similar $\bar{W}/\lambda \geq O(N)$ curve by assuming the user transmission radius can be increased to include $O(N^\alpha)$ other users, where α is between 0 and 1 and affects the delay tradeoff. This analysis does not consider the use of redundant packet transfers or multi-user reception. A similar approach by Toupis and Goldsmith in [16] shows that an improved tradeoff $\bar{W}/\lambda^2 = O(N \log^5(N))$ can be achieved when multi-user reception is used together with transmission radius scaling, but there was no proof of optimality.

In the context of a cell partitioned network as we have defined, an increased transmission radius would correspond to a user/cell density that is a function of N , that is, $d = O(N^\alpha)$. While our work was developed independently and intended only for the case $d = O(1)$ (independent of N), the necessary condition in Theorem 8 was proven for arbitrary values of the user/cell density d , and hence it can be used to evaluate the performance of the Toupis-Goldsmith algorithm applied to a cell partitioned network. Indeed, first note that the additional inequality $N\lambda \leq C$ must hold for any policy on a cell partitioned network (as the rate of new packets transmitted by their sources is less than or equal to C , the maximum number of transmissions possible during a slot). Thus, $\lambda \leq 1/d$ is necessary for any protocol, and directly plugging this inequality into (15) yields: $\bar{W}/\lambda^2 \geq \frac{(N-d)(1-\log(2))}{4}$.

Hence, the Toupis-Goldsmith algorithm is near-optimal over the class of all algorithms that can be implemented on a cell partitioned network that does not impose the constraint $d = O(1)$. We note that a recent preliminary result in [18] suggests that an improved tradeoff $\bar{W}/\lambda^3 \geq O(N)$ is possible if the network has different physical properties that allow for multi-hop transmission during a single slot (so that a bit can be transferred from node 1 to node 2, ..., to node K , all during a single slot). Of course, it is not possible to implement such an algorithm on the cell partitioned network that we have defined, because transmission on each successive hop would require a new timeslot.

VII. NON-I.I.D. MOBILITY MODELS

The analysis developed here for the i.i.d. mobility model can be used to bound the performance of a system with a Markovian mobility model. Instead of performing control actions on the network every slot, we decompose the network into a set of K parallel sub-networks. Packets are considered to be of ‘type- k ’ if they arrive during a timeslot t such that $t \bmod K = k$. On such timeslots, only control actions on type- k packets take place. The value of K is chosen suitably large to ensure that the user location distribution after K slots is within a constant factor of its steady state value. Specifically, if K is chosen such that, regardless of the initial configuration of users, the probability that two given users are in the same cell after K slots is at least $\frac{1}{2C}$, then delay under the three schemes is bounded by $O(KN)$, $O(K\sqrt{N})$, and $O(K \log(N))$, respectively. The $O(KN)$ result for the 2-hop relay algorithm (with no redundancy) follows by using the K -slot Lyapunov drift arguments developed in [1] [5]. The $O(K\sqrt{N})$, and $O(K \log(N))$ bounds follow by literally repeating the same arguments used for the \sqrt{N} redundancy algorithm and the Fair-Flooding algorithm on a K slot basis.

While this analysis offers a simple upper bound on average delay, we note that for many network models the value of K may depend on N , making these bounds larger than the $O(N)$ and $O(\sqrt{N})$ results for i.i.d. mobility. For example, the value of K for a Markovian random walk might be on the order of \sqrt{N} , representing the time required for a node to move from one side of the network to the other. However, it is possible that alternative scheduling schemes could yield lower delay. Indeed, in the next section it is shown through simulation that applying the 2-hop relay algorithm and the \sqrt{N} redundancy algorithm exactly as before (without the K -subchannel decomposition) yields similar performance for both i.i.d. and non-i.i.d. mobility. It may be possible to analytically establish this result by proving that the average re-visitation time between any two nodes remains $O(N)$ under Markovian random walks, and that the average time required to send out \sqrt{N} duplicates of a single packet to different nodes and the time required for the destination to encounter one of these duplicate-carrying users remains $O(\sqrt{N})$. We leave such questions for future work.

VIII. SIMULATION RESULTS

Here we compare the average delay obtained through both analysis and simulation as the network is scaled. We consider a network with cells given by an $M \times M$ grid as shown in Fig. 1. The number of cells C is equal to M^2 (where M is varied between 3 and 15 for simulations), and the number of users N is chosen as the even integer for which N/C most accurately approximates the optimal user/cell density value $d^* = 1.7933$.

In Fig. 5, plots of average end-to-end delay versus the number of users N are provided for the 2-hop relay algorithm and the $O(\sqrt{N})$ redundancy algorithm for both an i.i.d. and a non-i.i.d. mobility model. In the i.i.d. mobility model, users choose new cells uniformly over all cells in the network. In the non-i.i.d. model, each user chooses a new cell every timeslot according to the following Markovian dynamics: With

probability $\alpha < 1$ the user stays in the same cell, and else it moves to an adjacent cell to the North, South, East, or West, with each direction equally likely. In the case where a user is on the edge of the network and is selected to move in an infeasible direction, it stays in its place. Using standard random walk theory it is easy to verify that, in steady state, such a Markov model leaves users independently and uniformly distributed over all cells, as the stationary equation for the Markov chain is satisfied when all cell locations have equal probability [20] [21] [22]. In particular, if π_i represents the steady state probability of a particular cell i , we have:

$$\pi_i = \pi_i \alpha + \pi_a \frac{(1-\alpha)}{4} + \pi_b \frac{(1-\alpha)}{4} + \pi_c \frac{(1-\alpha)}{4} + \pi_d \frac{(1-\alpha)}{4}$$

where $\pi_a, \pi_b, \pi_c, \pi_d$ represent steady state probabilities for other cells, possibly including cell i . In the case when cell i is an interior cell, it has four distinct neighbors a, b, c, d . In the case when it is an edge cell with three neighbors a, b, c , we set $d = i$ (so that cell i is its own neighbor). In the case when cell i is a corner cell with 2 neighbors a and b , we set $c = d = i$. Clearly these steady state equations are satisfied when the π_i probabilities are set to $1/C$ for all i .³ Therefore, the network capacity μ is the same for both the i.i.d. mobility model and the non-i.i.d. mobility model, and is given by $\mu = \frac{p+q}{2d}$ as described in Theorem 1. In the simulation results we set the α parameter of the non-i.i.d. model to $\alpha = 1/2$.

For the capacity achieving 2-hop relay algorithm, the data rate λ into each user is fixed at 80% of the network capacity μ (given in Theorem 1), so that $\rho = \lambda/\mu = 0.8$. The top three curves for average delay in Fig. 5 respectively represent the exact analytical delay for i.i.d. mobility, the simulated performance of the i.i.d. mobility model, and the simulated performance of the Markovian mobility model. Note that the simulation curve for the i.i.d. mobility model is almost indistinguishable from the analytical curve $\mathbb{E}\{W\} = \frac{N-1-\lambda}{\mu-\lambda}$. The curves are plotted on a log log scale and have a slope of 1, indicating $O(N)$ delay. The delay curve for Markovian mobility is situated slightly above the curve for i.i.d. mobility, and also has a slope of 1. This suggests that for Markovian mobility, delay is increased by a constant multiplicative factor but remains $O(N)$.

Results for the \sqrt{N} redundancy protocol are also shown in the figure. Data rates λ are set to the value $\lambda = 0.8\tilde{\mu}$, where $\tilde{\mu}$ is given in (8). Note that, unlike the network capacity μ , the throughput $\tilde{\mu}$ decreases as $O(1/\sqrt{N})$. The analytical upper and lower bounds on delay for i.i.d. mobility are shown in the figure, each having a slope of $1/2$ indicating $O(\sqrt{N})$ growth (note that the lower bound represents the delay of sending just a single packet). The simulation performance for i.i.d. mobility is shown in the figure and is situated between the upper and lower bounds. The upper bound is larger than the simulated curve by approximately a factor of 10, suggesting that tighter bounds could be obtained through a more detailed analysis. The slope of the simulation curve varies between

³Similar results hold when the random walk has a different behavior at the edges. In particular, if the direction is chosen uniformly over all feasible directions, then the interior cells will have equal probability but the edge cells will have a different probability.




Fig. 5. Average delay versus the number of users N for the 2-hop relay algorithm and the \sqrt{N} redundancy algorithm.

$5/8$ and $1/2$. However, due to the $O(\sqrt{N})$ upper and lower bounds, the average slope would converge to $1/2$ if the graph were extended. Simulation of the Markovian mobility model is also provided, and the curve again lies slightly above the i.i.d. mobility curve. This suggests that delay under the Markovian model is close to $O(\sqrt{N})$.

Experiments to simulate the performance of the $O(\log(N))$ scheme were not performed. However, for this case, we would expect a discrepancy between the i.i.d. mobility model and the non-i.i.d. mobility model. Indeed, although the i.i.d. mobility model yields logarithmic delay, the delay under a Markovian mobility model would likely be closer to $O(\sqrt{N})$ due to the time required for a user to travel from one side of the network to the other.

IX. CONCLUSIONS

This work for the first time presents a multi-hop, multi-user system for which a relatively complete network theory can be developed. Exact expressions for network capacity were derived, and a fundamental rate-delay curve was established, representing performance bounds on throughput and end-to-end network delay for any conceivable routing and scheduling policy.

Delay analysis for the network was facilitated using a simple i.i.d. user mobility model. Under this model, an exact expression for end-to-end delay which includes the full effects of queueing was established for the capacity achieving 2-hop relay algorithm. Two other protocols which (necessarily) use redundant packet transfers were provided and shown to improve delay at the expense of reducing throughput. The rate-delay performance of these schemes was shown to lie on the boundary of the fundamental performance curve

$\text{delay}/\text{rate} \geq O(N)$. Analysis of general mobility models can be understood in terms of this i.i.d. analysis, where delay bounds can be scaled by the factor K , representing the number of slots required between sampling points for samples of user locations to look nearly i.i.d.. Furthermore, simulation results suggest that $O(\sqrt{N})$ delay can be achieved for networks with Markovian mobility, as the delay for such systems closely follows the delay curve for a system with i.i.d. mobility.

This inspires a rich set of questions concerning the fundamental limits of data networks. We believe that the condition $\text{delay}/\text{rate} \geq O(N)$ is necessary for general classes of mobile wireless networks, and that the $(\text{rate}, \text{delay}) = (O(1/\sqrt{N}), O(\sqrt{N}))$ operating point is always achievable. Such conjectures can perhaps be established using analytical techniques similar to those created here.

APPENDIX A — THE NETWORK CAPACITY THEOREM

Here we prove Theorem 1: *The capacity of a cell partitioned network is:*

$$\mu = \frac{p+q}{2d}$$

where p represents the probability of finding at least two users in a particular cell, and q represents the probability of finding a source-destination pair within a cell.

An algorithm for stabilizing the network whenever $\lambda < \mu$ is given in Section II-C. Here we prove $\lambda \leq \mu$ is necessary for stability.

Proof: (Necessity) Consider any stabilizing scheduling strategy, perhaps one that uses full knowledge of future events. Let $X_h(T)$ represent the total number of packets transferred over the network from sources to destinations in h hops during the interval $[0, T]$. Fix $\epsilon > 0$. For network stability, there must be arbitrarily large values T such that the sum output rate is within ϵ of the total input rate:

$$\frac{\sum_{h=1}^{\infty} X_h(T)}{T} \geq N\lambda - \epsilon \quad (25)$$

If this were not the case, the total number of packets in the network would grow to infinity and hence the network would be unstable. The total number of packet transmissions in the network during the first T slots is at least $\sum_{h=1}^{\infty} hX_h(T)$. This value must be less than or equal to the total number of transmission opportunities $Y(T)$, and hence:

$$\sum_{h=1}^{\infty} hX_h(T) \leq Y(T) \quad (26)$$

where $Y(T)$ represents the total number of cells containing at least 2 users in a particular timeslot, summed over all timeslots $1, 2, \dots, T$. By the law of large numbers, it is clear that $\frac{1}{T}Y(T) \rightarrow Cp$ as $T \rightarrow \infty$, where p is the steady state probability that there are two or more users within a particular cell, and is given by (2).

From (25) and (26), it follows that

$$\begin{aligned} \frac{1}{T}Y(T) &\geq \frac{1}{T}X_1(T) + \frac{2}{T} \sum_{h=2}^{\infty} X_h(T) \\ &\geq \frac{1}{T}X_1(T) + 2 \left((N\lambda - \epsilon) - \frac{1}{T}X_1(T) \right) \end{aligned}$$

and hence

$$\lambda \leq \frac{\frac{1}{T}Y(T) + \frac{1}{T}X_1(T) + 2\epsilon}{2N} \quad (27)$$

It follows that maximizing λ subject to (27) involves placing as much rate as possible on the single hop paths. However, the time average rate $\frac{1}{T}X_1(T)$ of 1-hop communication between source-destination pairs is bounded. Indeed, the probability q that a particular cell contains a source-destination pair during a timeslot can be written as 1 minus the probability that no such pair is present. For the source-destination matching $1 \leftrightarrow 2, 3 \leftrightarrow 4, \dots$, this probability is given as the value q specified in (3). Let $Z(T)$ represent the total number of cells containing source-destination pairs, summed over all timeslots $1, 2, \dots, T$. Again by the law of large numbers, it follows that $\frac{1}{T}Z(T) \rightarrow Cq$. Furthermore, it is clear that the number of packets delivered on one hop paths is less than or equal to the number of such opportunities:

$$\frac{1}{T}X_1(T) \leq \frac{1}{T}Z(T) \quad (28)$$

Combining constraints (27) and (28) and taking limits as $T \rightarrow \infty$, we have:

$$\lambda \leq \frac{Cp + Cq + 2\epsilon}{2N} \quad (29)$$

The necessary condition follows by using the user/cell density definition $d = N/C$, and noting that ϵ can be chosen to be arbitrarily small. \square

APPENDIX B — EXACT DELAY ANALYSIS OF THE 2-HOP RELAY ALGORITHM

Proof of Delay Bound in Theorem 3: A decoupled view of the network as perceived by a single user i is illustrated in Fig. 3. Due to the i.i.d. mobility, the source user can be represented as a Bernoulli/Bernoulli queue, where every timeslot a new packet arrives with probability λ , and a service opportunity arises with some fixed probability μ . We first show that $\mu = \frac{p+q}{2d}$. The Bernoulli nature of the server process implies that the transmission probability μ is equal to the time average rate of transmission opportunities of source i .⁴ Hence, we have $\mu = r_1 + r_2$, where r_1 represents the rate at which the source is scheduled to transmit directly to the destination, and r_2 represents the rate at which it is scheduled to transmit to one of its relay users. The cell partitioned relay algorithm schedules transmissions into and out of the relay nodes with equal probability, and hence r_2 is also equal to the rate at which the relay nodes are scheduled to transmit to the destination. The total rate of transmission opportunities over the network is thus $N(r_1 + 2r_2)$. A transmission opportunity occurs in any given cell with probability p , and hence:

$$Cp = N(r_1 + 2r_2) \quad (30)$$

⁴A transmission opportunity arises when a user is selected to transmit to another user, and corresponds to a service opportunity in the Bernoulli/Bernoulli queue. Such opportunities arise with probability μ every timeslot, independent of whether or not there is a packet waiting in the queue.

Recall that q is the probability that a given cell contains a source-destination pair. Since the cell partitioned relay algorithm schedules the single-hop ‘source-to-destination’ transmissions whenever possible, the rate r_1 satisfies:

$$Cq = Nr_1 \quad (31)$$

It follows from (31) that $r_1 = q/d$, and hence by (30) we infer that $r_2 = \frac{p-q}{2d}$. The total rate of transmissions out of the source node is thus given by $\mu = r_1 + r_2 = \frac{p+q}{2d}$.

The source is thus a Bernoulli/Bernoulli queue with input rate λ and server probability μ , having an expected number of packets given by $\bar{L}_{source} = \frac{\rho(1-\lambda)}{1-\rho}$, where $\rho \triangleq \lambda/\mu$ [23]. This queue is *reversible* ([22], [23]), and so the output process is also a Bernoulli stream of rate λ .

A given packet from this output process is transmitted to the *first relay node* with probability $\frac{r_2}{\mu(N-2)}$ (because with probability r_2/μ the packet is intended for a relay node, and each of the $N-2$ relay nodes are equally likely). Hence, every timeslot this relay independently receives a packet with probability $\tilde{\lambda} = \frac{\lambda r_2}{\mu(N-2)}$. The relay node is scheduled for a potential packet transmission to the destination with probability $\tilde{\mu} = \frac{r_2}{(N-2)}$ (because a ‘relay-to-destination’ opportunity arises in the relay node with probability r_2 , and arises for each of the $N-2$ destination nodes with equal probability). However, packet arrivals and transmission opportunities are mutually exclusive events in the relay node. It follows that the discrete time Markov Chain for queue occupancy in the relay node can be written as a simple birth-death chain which is identical to the chain of a continuous time M/M/1 queue with input rate $\tilde{\lambda}$ and service rate $\tilde{\mu}$ (where $\tilde{\lambda}/\tilde{\mu} = \rho$). This holds for each relay node, and the resulting occupancy at any relay is thus: $\bar{L}_{relay} = \frac{\rho}{1-\rho}$. From Little’s Theorem, the total network delay is $\bar{W}_i = [\bar{L}_{source} + (N-2)\bar{L}_{relay}]/\lambda$, and hence $\mathbb{E}\{W_i\} = \frac{N-1-\lambda_i}{\mu-\lambda_i}$, proving the theorem. \square

APPENDIX C — HETEROGENEOUS DATA RATES

Proof of Theorem 2: Here we prove that for heterogeneous data rates (λ_{ij}) such that there are at most K nonzero λ_{ij} entries in each row i , the symmetric capacity region satisfies:

$$\begin{aligned} \sum_j \lambda_{ij} &\leq \frac{(1 - e^{-d} - de^{-d})}{2d} + O(K/N) \quad \forall i \\ \sum_i \lambda_{ij} &\leq \frac{(1 - e^{-d} - de^{-d})}{2d} + O(K/N) \quad \forall j \end{aligned}$$

Before proving the theorem, we first note that whenever $N > d$, we have:

$$e^{\frac{-d^2}{N-d}} e^{-d} \leq \left(1 - \frac{d}{N}\right)^N \leq e^{-d}$$

which can be proven by taking the logarithm of the above inequality and using the fact that $\log(1+x) \leq x$ whenever $x > -1$.⁵ The difference between the upper and lower bounds is thus $e^{-d} \left(1 - e^{\frac{-d^2}{N-d}}\right)$. Using the Taylor expansion $e^{\frac{-d^2}{N-d}} =$

$1 + \frac{-d^2}{N-d} + O(1/N^2)$ reveals that this difference is $O(1/N)$, and hence $(1 - \frac{d}{N})^N = e^{-d} + O(1/N)$.

Proof: (Necessity) The proof that the above inequalities are necessary conditions for stability is similar to the proof of Theorem 1, where the equation (25) is replaced by:

$$\frac{1}{T} \sum_{h=1}^{\infty} X_h(T) \geq \sum_i \sum_j \lambda_{ij} - \epsilon$$

Repeating the same argument as in Theorem 1, it follows that [compare with (29)]:

$$\frac{1}{N} \sum_i \sum_j \lambda_{ij} \leq \frac{Cp + C\tilde{q} + 2\epsilon}{2N} = \frac{p}{2d} + \frac{\tilde{q}}{2d} + \frac{\epsilon}{N}$$

where p is the probability that at least two users are within a cell (given in (2)), and \tilde{q} is the probability that there exists a source-destination pair within the cell. Note that \tilde{q} may be different from the value of q given in (3) because of the different sets of source-destination pairs. However, since each user i has at most K destination nodes to consider, the union bound implies that the probability of any particular user entering a given cell along with at least one of its destinations is less than or equal to $\frac{1}{C} \frac{K}{C}$, so that $\tilde{q} \leq \frac{N}{C} \frac{K}{C} = O(K/N)$.

The probability p that at least two users are within a cell satisfies:

$$\begin{aligned} p &= 1 - \left(1 - \frac{d}{N}\right)^N - d \left(1 - \frac{d}{N}\right)^{N-1} \\ &= 1 - e^{-d} - de^{-d} + O(1/N) \end{aligned}$$

Hence, $\frac{1}{N} \sum_i \sum_j \lambda_{ij} \leq \frac{1 - e^{-d} - de^{-d}}{2d} + O(K/N)$. This together with the fact that no user sends or receives more than any other proves the result. \square

For sufficiency, we consider a 2-hop routing scheme, where data is routed uniformly over all relay nodes on the first hop regardless of its destination. We note that such a *traffic uniformization* scheme is conceptually similar to the 2-stage switch scheduling algorithm developed for $N \times N$ packet switches in [24], where packets are randomly assigned to output ports at the first stage so that traffic is uniform at the second stage.

Proof: (Sufficiency) From the Network Capacity Theorem developed in [1] [5], we know that it is sufficient to describe a transmission strategy yielding long term node-to-node packet exchange rates μ_{ij} together with a set of multi-commodity flows which route all data to their destinations without exceeding these rates on any link (i, j) . Consider the strategy of choosing a transmitter and receiver in each cell completely randomly over all user pairs. As the expected number of packet transfer opportunities over the network is Cp opportunities per slot, the total rate of opportunities between any two links is $\mu_{ij} = \frac{Cp}{N(N-1)}$.

Suppose now the rate of exogenous data arriving to any node i is identically λ (for some data rate λ), as is the sum rate of data entering the network destined for any node j , so that $\sum_i \lambda_{ij} = \sum_j \lambda_{ij} = \lambda$ for all i, j . (Any smaller rate matrix which does not sum to λ in every row and column can be increased to a matrix which does have this property). Consider the 2-hop routing scheme were exogenous packets at

⁵Note that $\frac{-d}{N-d} \leq -\log\left(1 + \frac{d}{N-d}\right) = \log\left(1 - \frac{d}{N}\right) \leq \frac{-d}{N}$.

a source are routed randomly and uniformly to any available relay node, and these packets are then transferred from relay to destination. Since on the first hop the algorithm routes data independently of its destination, the incoming traffic to the relay nodes is uniformly distributed, so that each relay receives data destined for node j at a rate $\lambda/(N-1)$ for all destinations j .

The total rate of traffic flowing over any link from i to j is thus $2\lambda/(N-1)$ (where a stream of total rate $\lambda/(N-1)$ flows from i to j due to packets from source i being relayed to j , and data of rate $\lambda/(N-1)$ flows from i to j due to traffic being relayed from i to destination j). This traffic satisfies the link constraint provided that $2\lambda/(N-1) \leq \mu_{ij} = \frac{Cp}{N(N-1)}$, or equivalently that $\lambda \leq \frac{p}{2d}$. Thus, any rate matrices (λ_{ij}) satisfying $\sum_i \lambda_{ij} \leq \frac{p}{2d}$ for all j and $\sum_j \lambda_{ij} \leq \frac{p}{2d}$ for all i are within the capacity region, where $\frac{p}{2d} = \frac{1-e^{-d}-de^{-d}}{2d} + O(1/N)$. \square

APPENDIX D — MINIMUM DELAY FOR 2-HOP ROUTING

Here we derive a recursive formula for the minimum average delay for sending a single packet from source to destination in the case when routing is restricted to 2-hop paths. We assume that multi-user reception is not available, so that at most one user per cell can receive a packet during a single timeslot.

The minimum delay algorithm transfers the packet to its destination whenever the source or a duplicate-carrying relay is in the same cell as the destination, and otherwise schedules the source to deliver a duplicate version of the packet to a new user whenever possible. Let $\mathbb{E}\{T_N\}$ represent the expected time for the packet to reach the destination. The value of $\mathbb{E}\{T_N\}$ can be computed recursively by defining variables X_1, X_2, \dots, X_{N-1} , where X_k represents the expected time for the packet to reach its destination given that k users are carrying duplicates of the packet. The probability that a particular user does not move to the same cell as the destination during a timeslot is $(1 - 1/C)$. Therefore, the probability that at least one user among a group of k users *does* reach the destination is $1 - (1 - 1/C)^k$. Note that because all paths are restricted to 2 hops, the number of users holding a duplicate version of the packet increases by at most one every slot. This number stays the same if the source user does not visit anyone new, and if (independently) all $k-1$ other users holding the packet do not visit the destination. Considering the Markov nature of the problem, we have the following transition probabilities for each state $k \in \{1, \dots, N-2\}$:

$$\begin{aligned} Pr[k \rightarrow \text{end}] &= 1 - \left(1 - \frac{1}{C}\right)^k \\ Pr[k \rightarrow k] &= \left(1 - \frac{1}{C}\right)^{N-k} \left(1 - \frac{1}{C}\right)^{k-1} \\ &= \left(1 - \frac{1}{C}\right)^{N-1} \\ Pr[k \rightarrow k+1] &= 1 - Pr[k \rightarrow \text{end}] - Pr[k \rightarrow k] \end{aligned}$$

In state $k = N-1$, the remaining time to finish is a geometric variable with probability $1 - (1 - \frac{1}{C})^{N-1}$. The

values of X_i can thus be computed recursively as follows:

$$\begin{aligned} X_{N-1} &= \frac{1}{1 - (1 - 1/C)^{N-1}} \\ X_k &= 1 + X_k(1 - 1/C)^{N-1} + \\ &\quad X_{k+1} [(1 - 1/C)^k - (1 - 1/C)^{N-1}] \end{aligned}$$

and $\mathbb{E}\{T_N\} = X_1$.

APPENDIX E — MULTI-USER RECEPTION

Here we show that multi-user reception cannot overcome the \sqrt{N} lower bound on delay for 2-hop routing. Specifically, we show that the delay $\mathbb{E}\{T_N\}$ for any algorithm which restricts packets to 2-hop paths satisfies:

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}\{T_N\}}{\sqrt{N}} \geq e^{-d^2}$$

Proof: Consider sending a single packet to its destination over an empty network. Let K_t represent the total number of users who have the packet at the beginning of slot t (not including the destination). Since scheduling restricts transfers to 2-hop paths, the number of users holding the packet increases every timeslot by at most the number of users in the same cell as the source (which is $d - 2/C$ on average). Hence, we have for all $t \geq 1$:

$$\mathbb{E}\{K_t\} \leq td \quad (32)$$

Note that during slots $\{1, 2, \dots, t\}$ there are at most K_t users holding the packet, and hence during each of these slots the probability that no packet-holding user enters the cell of the destination is at least $(1 - \frac{1}{C})^{K_t}$. Thus:

$$\begin{aligned} Pr[T_N > t | K_t] &\geq \left(1 - \frac{1}{C}\right)^{tK_t} \\ &= \left(1 - \frac{d}{N}\right)^{tK_t} \end{aligned} \quad (33)$$

We thus have:

$$\mathbb{E}\{T_N\} \geq tPr[T_N > t] \quad (34)$$

$$= t\mathbb{E}_{K_t}\{Pr[T_N > t | K_t]\} \quad (35)$$

$$\geq t\mathbb{E}_{K_t}\left\{\left(1 - \frac{d}{N}\right)^{tK_t}\right\} \quad (36)$$

$$\geq t\left(1 - \frac{d}{N}\right)^{t\mathbb{E}\{K_t\}} \quad (37)$$

$$\geq t\left(1 - \frac{d}{N}\right)^{t^2 d} \quad (38)$$

where inequality (36) follows from (33), inequality (37) holds by Jensen's inequality (noticing that the function β^x is convex in x for any $\beta > 0$), and (38) follows from (32). This holds for all integers t . Choosing $t = \sqrt{N}$ yields $\mathbb{E}\{T_N\} \geq \sqrt{N} \left(1 - \frac{d}{N}\right)^{Nd} \rightarrow e^{-d^2} \sqrt{N}$. \square

APPENDIX F — DELAY OF \sqrt{N} REDUNDANCY ALGORITHM

Here we prove eq. (9), establishing an $O(\sqrt{N})$ bound on the service time $\mathbb{E}\{T_N\}$ for the partial feedback scheme with \sqrt{N} redundancy. The proof requires the following preliminary lemma.

Lemma 5: Consider N users which independently choose to enter one of C cells, and recall that $d = N/C$ represents the expected number of users per cell. Let J represent the number of users contained in a given cell. We have:⁶

$$\mathbb{E}\{J | J \geq 1\} \leq 1 + d$$

Proof: Let I_i represent an indicator variable taking the value 1 if the i th user of the subset is in the cell, and 0 otherwise. Define K as the lowest indexed user within the cell, where we let $K = N + 1$ if no users are present. Thus, $J = \sum_{i=K}^N I_i$. We have:

$$\begin{aligned} \mathbb{E}\{J | J \geq 1\} &= 1 + \mathbb{E}\left\{\sum_{i=K+1}^N I_i \mid J \geq 1\right\} \\ &= 1 + \mathbb{E}_{K|J \geq 1}\left\{\sum_{i=K+1}^N \mathbb{E}\{I_i | K, J \geq 1\} \mid J \geq 1\right\} \\ &= 1 + \mathbb{E}_{K|J \geq 1}\left\{\sum_{i=K+1}^N \frac{1}{C} \mid J \geq 1\right\} \quad (39) \\ &\leq 1 + \mathbb{E}_{K|J \geq 1}\left\{\sum_{i=1}^N \frac{1}{C} \mid J \geq 1\right\} = 1 + \frac{N}{C} \end{aligned}$$

where (39) follows because the condition $J \geq 1$ can be inferred by knowledge of K , and $\mathbb{E}\{I_i | K\} = \frac{1}{C}$ for all $i > K$. Indeed, the event $K = k$ is equivalent to the event that user k is in the cell but users $1, \dots, k-1$ are not in the cell, and this event is independent of the location of users $i \in \{k+1, \dots, N\}$. \square

To prove the \sqrt{N} bound on $\mathbb{E}\{T_N\}$, recall that $T_N = S_1 + S_2$, where S_1 represents the time required for the source to send out \sqrt{N} replicas of the packet (while competing with other sessions for network resources), and S_2 represents the time required to reach the destination given that \sqrt{N} users have the packet.

Lemma 6:

$$\mathbb{E}\{S_1\}, \mathbb{E}\{S_2\} \leq \frac{4+2d}{\gamma_N(1-e^{-d})} \sqrt{N}$$

where γ_N is a sequence that converges to 1 as $N \rightarrow \infty$.

Proof: The $\mathbb{E}\{S_1\}$ bound: Let S_1 represent the time required for the source to deliver a duplicate packet to \sqrt{N} distinct users. For the duration of S_1 , there are at least $N - \sqrt{N}$ users who do not have the packet, and hence every timeslot the probability that at least one of these users visits the cell of the source is at least $1 - (1 - \frac{1}{C})^{N-\sqrt{N}}$. Given this event, the probability that the source is chosen by the partial feedback algorithm to transmit is expressed by the product $\alpha_1 \alpha_2$, representing probabilities for the following conditionally independent events: α_1 is the probability that the source is selected from all other users in the cell to be the transmitting user, and α_2 represents the probability that this source is chosen to operate in ‘source-to-relay’ mode. Let random variable J represent the number of additional users in

the cell of the source (excluding the source user itself). The value of α_1 is thus $\alpha_1 = \mathbb{E}\{1/(J+1) | J \geq 1\}$. By Jensen’s inequality, we have:

$$\begin{aligned} \alpha_1 &\geq 1/\mathbb{E}\{1+J | J \geq 1\} \\ &\geq 1/(2+d) \end{aligned}$$

where the last inequality follows because $\mathbb{E}\{J | J \geq 1\} \leq 1 + d$ (as proven in Lemma 5).

The probability α_2 that the source operates in ‘source-to-relay’ mode is $1/2$. Thus, every timeslot during the interval S_1 , the source delivers a replica packet to a new user with probability of at least ϕ , where

$$\begin{aligned} \phi &\geq \left(1 - \left(1 - \frac{1}{C}\right)^{N-\sqrt{N}}\right) \frac{1}{2(2+d)} \\ &\rightarrow \frac{1 - e^{-d}}{4 + 2d} \end{aligned}$$

The average time until a replica is transmitted to a new user is thus a geometric variable with mean less than or equal to $1/\phi$. It is possible that two or more replicas are delivered in a single timeslot. However, in the worst case, \sqrt{N} of these times are required, so that the average time $\mathbb{E}\{S_1\}$ is upper bounded by \sqrt{N}/ϕ . \square

Proof: The $\mathbb{E}\{S_2\}$ bound: To prove the bound on $\mathbb{E}\{S_2\}$, note that every timeslot in which there are at least \sqrt{N} users with replicas of the packet, the probability that one of these users transmits the packet to the destination is given by the chain of probabilities $\theta_0 \theta_1 \theta_2 \theta_3$. The θ_i values represent probabilities for the following conditionally independent events: θ_0 represents the probability that there is at least one other user in the same cell as the destination (and is given by $\theta_0 = 1 - (1 - 1/C)^{N-1} \rightarrow 1 - e^{-d}$), θ_1 represents the probability that the destination is selected as the receiver (where, similar to the α_1 computation, we have $\theta_1 \geq 1/(2+d)$), θ_2 represents the probability that the sender operates in ‘relay-to-destination’ mode (where $\theta_2 = 1/2$), and θ_3 represents the probability that the sender is one of the \sqrt{N} users who have a replica of the packet intended for the destination (where $\theta_3 = \sqrt{N}/(N-1) \geq 1/\sqrt{N}$). Thus, every timeslot, the probability that the S_2 time comes to completion is at least $\frac{(1-e^{-d})}{(4+2d)\sqrt{N}}$. The value of $\mathbb{E}\{S_2\}$ is thus less than or equal to the inverse of this quantity. \square

APPENDIX G — LOGARITHMIC DELAY FOR FLOODING PROTOCOL

Here we prove Lemma 3: Under the algorithm of flooding the network with a single packet, for any network size $N \geq \max\{d, 2\}$, the expected time $\mathbb{E}\{T_N\}$ for the packet to reach every user satisfies $\mathbb{E}\{T_N\} \leq \mathbb{E}\{S_1\} + \mathbb{E}\{S_2\}$, where:

$$\begin{aligned} \mathbb{E}\{S_1\} &\leq \frac{\log(N)(1+d/2)}{\log(2)(1-e^{-d/2})} \\ \mathbb{E}\{S_2\} &\leq 1 + \frac{2}{d}(1 + \log(N/2)) \end{aligned}$$

Proof: (The $\mathbb{E}\{S_2\}$ Bound) Let M represent the number of users who do not initially have the packet (so that $M \leq N/2$),

⁶An exact value of $\mathbb{E}\{J | J \geq 1\} = \mathbb{E}\{J\}/Pr[J \geq 1]$ can easily be computed and leads to tighter but more complicated delay bounds.

and label these M users $\{u_1, u_2, \dots, u_M\}$. Let X_i represent the number of timeslots it takes for the non-packet holding user u_i to reach a cell containing a user who possesses a packet. Due to the multi-user reception feature, user u_i must receive the packet at this time. The random variable X_i is geometric, in that a 'success' happens on any given timeslot with probability $\psi \geq 1 - (1 - \frac{1}{C})^{N/2}$. Thus, we have for all N :

$$\psi \geq 1 - e^{-d/2} \quad (40)$$

All times X_i are independent and identically distributed, and hence the random variable S_2 is equal to the maximum value of at most $M = \lfloor N/2 \rfloor$ i.i.d. variables. Hence, $\mathbb{E}\{S_2\} \leq \mathbb{E}\{\max\{X_1, X_2, \dots, X_M\}\}$. To obtain a simple bound on this time, we consider new random variables $\{Y_1, Y_2, \dots, Y_M\}$ which are i.i.d. and *exponentially distributed* with rate $\lambda = \log(1/(1 - \psi))$. Notice that the random variable $1 + Y_i$ is *stochastically greater* than X_i , because the complementary distribution functions satisfy $\Pr[1 + Y_i > t] \geq \Pr[X_i > t]$ for all real numbers t (see [19] for a treatment of stochastic dominance for random variables). It follows that:

$$\begin{aligned} \mathbb{E}\{S_2\} &\leq \mathbb{E}\{\max\{X_1, X_2, \dots, X_M\}\} \\ &\leq 1 + \mathbb{E}\{\max\{Y_1, Y_2, \dots, Y_M\}\} \end{aligned}$$

The expected maximum of M i.i.d. exponential variables of rate λ is equal to the expectation of the sum of intervals $I_1 + I_2 + \dots + I_M$, where I_i represents the duration of time between the $(i - 1)^{th}$ and i^{th} completion time. The interval I_1 is the first completion time of M independently racing exponential variables, and hence I_1 is exponentially distributed with rate $M\lambda$. Furthermore, I_2 is the first completion time of $M - 1$ racing exponential variables, I_3 is the first completion time of $M - 2$ racing exponentials, and so on. It follows that:

$$\mathbb{E}\{I_1 + I_2 + \dots + I_M\} = \frac{1}{\lambda} \sum_{m=1}^M \frac{1}{m}$$

Hence, $\mathbb{E}\{S_2\} \leq 1 + \frac{1}{\lambda} \sum_{m=1}^M \frac{1}{m}$, which is upper bounded by $1 + \frac{1}{\lambda}(1 + \log(M))$. Hence:

$$\mathbb{E}\{S_2\} \leq 1 + \frac{1 + \log(M)}{\log(1/(1 - \psi))} \leq 1 + \frac{1 + \log(N/2)}{\log(e^{d/2})}$$

□

Proof: (The $\mathbb{E}\{S_1\}$ bound) We compute a bound on $\mathbb{E}\{S_1\}$ by noting that $\mathbb{E}\{S_1\} \leq \mathbb{E}\{\tilde{S}_1\}$, where \tilde{S}_1 is the time to reach at least $N/2$ users when the multi-user reception feature is turned off, and any transmitted packet is received by at most 1 other user within a cell. It turns out that the variable \tilde{S}_1 is easier to work with, as the number of users holding the packet can at most double every timeslot. Let K_t represent the number of users containing a duplicate version of the packet at timeslot $t \in 1, 2, \dots$ (suppose only the source user has the packet at time 0, so that $K_0 = 1$). Let u_1, u_2, \dots, u_{K_t} represent the users containing the packet at time t . Each of these users u_i delivers the packet to a_i new users on the next timeslot, where a_i is a binary random variable taking a value of either 0 or 1. Whenever there are at least $N/2$ users which do not currently hold the packet, we have that $\mathbb{E}\{a_i\} \geq \theta_1 \theta_2$,

where $\theta_1 = 1 - (1 - \frac{1}{C})^{N/2}$ represents a lower bound on the probability that at least one of the new users enters the cell of user u_i , and θ_2 represents a lower bound on the probability that user i is selected to transmit *its* replica among all other packet-holding users within the cell. Define J as the total number of other packet-holding users in the cell (not including user i). It follows that:

$$\theta_2 = \mathbb{E}\left\{\frac{1}{1 + J}\right\} \quad (41)$$

$$\geq \frac{1}{1 + \mathbb{E}\{J\}} \quad (42)$$

$$\geq \frac{1}{1 + d/2} \quad (43)$$

where (42) follows by Jensen's inequality and convexity of the function $1/(1+x)$, and (43) follows because there are no more than $N/2$ packet holding users, and hence $\mathbb{E}\{J\} \leq \frac{N}{2C} = d/2$. Thus:

$$\begin{aligned} \mathbb{E}\{a_i\} &\geq \frac{1 - (1 - \frac{1}{C})^{N/2}}{1 + d/2} \\ &\geq \frac{1 - e^{-d/2}}{1 + d/2} \end{aligned} \quad (44)$$

where (44) follows because $(1 - \frac{d}{N})^N \leq e^{-d}$ for all $N \geq d > 0$.

Let $Z_t = K_t/K_{t-1}$ be a random variable representing the multiplicative factor by which the number of packet-holding users grows after one timeslot. (Note that $1 \leq Z_t \leq 2$). It clearly holds that:

$$Z_{t+1} = \frac{K_t + a_1 + a_2 + \dots + a_{K_t}}{K_t}$$

The a_i random variables are not independent, although they are identical. Thus, for any timeslot t in which fewer than $N/2$ users have packets:

$$\begin{aligned} \mathbb{E}\{Z_{t+1} | K_t\} &= \frac{K_t + K_t \mathbb{E}\{a_1\}}{K_t} \\ &= 1 + \mathbb{E}\{a_1\} \\ &\geq 1 + \frac{1 - e^{-d/2}}{1 + d/2} \end{aligned} \quad (45)$$

Now consider the stopping time \tilde{S}_1 where at $t = \tilde{S}_1 - 1$ there are fewer than $N/2$ users with packets, but at time $t = \tilde{S}_1$ the $N/2$ threshold is either met or crossed. Note that \tilde{S}_1 is similar to a *stopping time variable*, treated in [22], [19], although the event $\{\tilde{S}_1 \geq t\}$ is not independent of Z_t . The number of users $K_{\tilde{S}_1}$ containing the packet at time $t = \tilde{S}_1$ satisfies:

$$N \geq K_{\tilde{S}_1} = Z_1 Z_2 \dots Z_{\tilde{S}_1}$$

and hence

$$\log(N) \geq \log(Z_1) + \log(Z_2) + \dots + \log(Z_{\tilde{S}_1})$$

Define the indicator random variable I_t to be 1 if $\tilde{S}_1 \geq t$, and 0 otherwise. Taking expectations of the above inequality, we find:

$$\begin{aligned}
 \log(N) &\geq \mathbb{E} \left\{ \sum_{t=1}^{\tilde{S}_1} \log(Z_t) \right\} \\
 &= \mathbb{E} \left\{ \sum_{t=1}^{\infty} \log(Z_t) I_t \right\} \\
 &= \mathbb{E} \left\{ \sum_{t=1}^{\infty} \mathbb{E} \{ \log(Z_t) I_t \mid K_{t-1} \} \right\} \\
 &= \mathbb{E} \left\{ \sum_{t=1}^{\infty} I_t \mathbb{E} \{ \log(Z_t) \mid K_{t-1} \} \right\}
 \end{aligned}$$

where the last equality follows because the variable K_{t-1} completely determines the binary value of I_t . Recall that $1 \leq Z_t \leq 2$, and hence $\log(Z_t) \geq \log(2)(Z_t - 1)$ (as the lower bound values are points along the chord of the concave function $\log(Z)$ over the interval $1 \leq Z \leq 2$). We thus have:

$$\begin{aligned}
 \frac{\log(N)}{\log(2)} &\geq \mathbb{E} \left\{ \sum_{t=1}^{\infty} I_t \mathbb{E} \{ (Z_t - 1) \mid K_{t-1} \} \right\} \\
 &\geq \left(\frac{1 - e^{-d/2}}{1 + d/2} \right) \mathbb{E} \left\{ \sum_{t=1}^{\infty} I_t \right\} \\
 &= \left(\frac{1 - e^{-d/2}}{1 + d/2} \right) \mathbb{E} \{ \tilde{S}_1 \}
 \end{aligned} \tag{46}$$

where (46) follows from (45). Thus, $\mathbb{E} \{ \tilde{S}_1 \} \leq \mathbb{E} \{ S_1 \} \leq \frac{\log(N)(1+d/2)}{\log(2)(1-e^{-d/2})}$. \square

APPENDIX H — MINIMUM DELAY FOR MULTI-HOP ROUTING IS LOGARITHMIC

Lemma 7: Starting with a single packet contained in one user in an empty network of size N , the flooding algorithm of delivering the packet to its destination by having every duplicate-carrying user transmit to other users whenever possible has an average delay $\mathbb{E} \{ T_N \}$ which is logarithmic. In particular

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E} \{ T_N \}}{\log(N)} \geq \frac{1}{\log(1+d)}$$

This bound holds even if multi-user reception is available.

Proof: As in the proof of Lemma 3, define K_t as the number of users holding the packet at time t (where $K_0 = 1$), and let $Z_t = K_t/K_{t-1}$ represent the growth factor after one timeslot. We have:

$$Z_{t+1} = \frac{K_t + a_1 + a_2 + \dots + a_{K_t}}{K_t}$$

where a_i represents the number of new users to which the i th packet-holding user transmits during a timeslot. We clearly have $\mathbb{E} \{ a_i \} \leq d$ during any timeslot, and hence:

$$\mathbb{E} \{ Z_{t+1} \mid K_t \} = \frac{K_t + K_t \mathbb{E} \{ a_1 \}}{K_t} \leq 1 + d$$

Because $K_t = Z_1 Z_2 \dots Z_t$, it follows by recursion that:

$$\mathbb{E} \{ K_t \} \leq (1+d)^t \tag{47}$$

Note that during slots $\{1, 2, \dots, t\}$ there are at most K_t users holding the packet, so the probability that none of these users enters the cell of the destination on such a timeslot is greater than or equal to $(1 - \frac{1}{C})^{K_t}$. Hence, the proof given in Appendix E for the \sqrt{N} bound for 2-hop routing can be followed exactly up to (37). In particular, we have [compare with (33)-(37)]:

$$\begin{aligned}
 \mathbb{E} \{ T_N \} &\geq t \Pr \{ T_N > t \} \\
 &= t \mathbb{E}_{K_t} \{ \Pr \{ T_N > t \mid K_t \} \} \\
 &\geq t \mathbb{E}_{K_t} \left\{ \left(1 - \frac{d}{N} \right)^{t K_t} \right\} \\
 &\geq t \left(1 - \frac{d}{N} \right)^{t \mathbb{E} \{ K_t \}}
 \end{aligned}$$

Using (47) in the above inequality, we have:

$$\mathbb{E} \{ T_N \} \geq t \left(1 - \frac{d}{N} \right)^{t(1+d)^t} \tag{48}$$

The above inequality holds for all integers $t \geq 0$. For convenience, we choose t to represent a base $(1+d)$ logarithm: $t \triangleq \log_{1+d}(\alpha N^\beta)$, where β is any number less than 1, and α is chosen within the bounds $1 \leq \alpha \leq (d+1)$ so that t is an integer. Using this value of t in (48), we have:

$$\mathbb{E} \{ T_N \} \geq \frac{(\log(\alpha) + \beta \log(N))}{\log(1+d)} \left[\left(1 - \frac{d}{N} \right)^N \right]^{\frac{\alpha N^\beta \log(\alpha N^\beta)}{N \log(1+d)}}$$

Note that $(1 - \frac{d}{N})^N \rightarrow e^{-d}$ as $N \rightarrow \infty$, and its exponent $\frac{\alpha N^\beta \log(\alpha N^\beta)}{N \log(1+d)}$ converges to 0 whenever $\beta < 1$. It follows that

$$\left[\left(1 - \frac{d}{N} \right)^N \right]^{\frac{\alpha N^\beta \log(\alpha N^\beta)}{N \log(1+d)}} \rightarrow 1, \text{ and hence:}$$

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E} \{ T_N \}}{\log(N)} \geq \frac{\beta}{\log(1+d)}$$

for any $\beta < 1$. The bound can be optimized by taking a limit as $\beta \rightarrow 1$, yielding the result. \square

APPENDIX I — TAKING INFIMUMS OVER SETS

Here we compute $\inf_{\Theta} \mathbb{E} \{ X \mid \Theta \}$ for a non-negative random variable X , where the infimum is taken over all events Θ such that $\Pr[\Theta] \geq \frac{1}{2}$. Let $P(x) = \Pr[X \leq x]$ represent the cumulative distribution function for X . Let ω be the unique real number such that $\Pr[X < \omega] \leq \frac{1}{2}$ and $\Pr[X \leq \omega] \geq \frac{1}{2}$. Note that if $P(x)$ is continuous, then $\Pr[X < \omega] = \Pr[X \leq \omega] = \frac{1}{2}$. In general, a non-continuous distribution may have a point mass at $x = \omega$.

Lemma 8: For any non-negative random variable X , we have:

$$\begin{aligned}
 \inf_{\{\Theta \mid \Pr[\Theta] \geq \frac{1}{2}\}} \mathbb{E} \{ X \mid \Theta \} &= \mathbb{E} \{ X \mid X < \omega \} 2\Pr[X < \omega] + \\
 &\quad \omega (1 - 2\Pr[X < \omega])
 \end{aligned}$$

Note that the infimum depends only on the cumulative distribution function $P(x)$. In the special case when $P(x)$ is continuous at $x = \omega$, then $\Pr[X < \omega] = \Pr[X \leq \omega] = \frac{1}{2}$,

and hence the lemma implies that the infimum is equal to $\mathbb{E}\{X \mid X \leq \omega\}$.

Proof: To prove the lemma, let $p(x) \triangleq \frac{dP(x)}{dx}$ represent the generalized density function of X (which may contain impulses if $P(x)$ is not continuous). Consider any event Θ such that $Pr[\Theta] \geq \frac{1}{2}$. Define the conditional probability distribution $f(x) \triangleq p_{X|\Theta}(x|\Theta)$. Note that $p(x) = p_{X|\Theta}(x|\Theta)Pr[\Theta] + p_{X|\Theta^c}(x|\Theta^c)Pr[\Theta^c]$ (where Θ^c represents the complement of the event Θ). Hence, $p_{X|\Theta}(x|\Theta) \leq p(x)/Pr[\Theta] \leq p(x)/\frac{1}{2}$. That is:

$$f(x) \leq 2p(x) \quad \text{for all } x \quad (49)$$

Note also that $f(x)$ is a probability distribution for a non-negative variable, so that $\int_0^\infty f(x)dx = 1$. We have:

$$\begin{aligned} \mathbb{E}\{X \mid \Theta\} &= \int_0^{\omega^-} xf(x)dx + \int_{\omega^-}^\infty xf(x)dx \\ &= \int_0^{\omega^-} x2p(x)dx + \int_0^{\omega^-} x[f(x) - 2p(x)]dx \\ &\quad + \int_{\omega^-}^\infty xf(x)dx \\ &\geq \int_0^{\omega^-} x2p(x)dx + \omega \int_0^{\omega^-} [f(x) - 2p(x)]dx \\ &\quad + \omega \int_{\omega^-}^\infty f(x)dx \end{aligned} \quad (50)$$

where (50) follows because (49) implies the integrand of the second integral is non-positive for all x (so that $\int_0^{\omega^-} x[f(x) - 2p(x)]dx \geq \omega \int_0^{\omega^-} [f(x) - 2p(x)]dx$). Noting that $\int_0^{\omega^-} f(x)dx + \int_{\omega^-}^\infty f(x)dx = 1$, inequality (50) implies:

$$\begin{aligned} \mathbb{E}\{X \mid \Theta\} &\geq \mathbb{E}\{X \mid X < \omega\} 2Pr[X < \omega] + \\ &\quad \omega - \omega \int_0^{\omega^-} 2p(x)dx \\ &= \mathbb{E}\{X \mid X < \omega\} 2Pr[X < \omega] + \\ &\quad \omega (1 - 2Pr[X < \omega]) \end{aligned} \quad (51)$$

The lower bound (51) holds for all events Θ such that $Pr[\Theta] \geq 1/2$, and hence:

$$\inf_{\{\Theta \mid Pr[\Theta] \geq \frac{1}{2}\}} \mathbb{E}\{X \mid \Theta\} \geq \mathbb{E}\{X \mid X < \omega\} 2Pr[X < \omega] + \omega (1 - 2Pr[X < \omega])$$

We now show that the reverse inequality is also true. Let A be the outcome of a biased coin flip that is independent of X . Specifically, let $Pr[A = 1] = q$, $Pr[A = 0] = 1 - q$, where q is the value such that $qPr[X = \omega] = (\frac{1}{2} - Pr[X < \omega])$. Note that $0 \leq q \leq 1$ because $Pr[X = \omega] + Pr[X < \omega] \geq \frac{1}{2}$ but $Pr[X < \omega] \leq \frac{1}{2}$.

Consider the particular event Θ^* defined as follows:

$$\Theta^* \triangleq \{\{X < \omega\} \cup \{\{X = \omega\} \cap \{A = 1\}\}\} \quad (52)$$

That is, Θ^* represents the event that either $X < \omega$, or both $X = \omega$ and $A = 1$. Note that $Pr[\Theta^*] = 1/2$, because

$Pr[\Theta^*] = Pr[X < \omega] + qPr[X = \omega]$. We then have:

$$\begin{aligned} \mathbb{E}\{X \mid \Theta^*\} &= \mathbb{E}\{X \mid X < \omega\} \frac{Pr[X < \omega]}{Pr[\Theta^*]} + \omega \frac{qPr[X = \omega]}{Pr[\Theta^*]} \\ &= \mathbb{E}\{X \mid X < \omega\} 2Pr[X < \omega] + \\ &\quad \omega (1 - 2Pr[X < \omega]) \end{aligned}$$

Thus, the particular event Θ^* allows the conditional expectation to meet the lower bound of (51). Thus, Θ^* is the minimizing event, and its resulting expectation is equal to the infimum, proving the lemma. \square

REFERENCES

- [1] M. J. Neely. *Dynamic Power Allocation and Routing for Satellite and Wireless Networks with Time Varying Channels*. PhD thesis, Massachusetts Institute of Technology, LIDS, 2003.
- [2] M. J. Neely and E. Modiano. Capacity and delay tradeoffs for ad-hoc mobile networks. *(Invited Paper) IEEE BroadNets, San Jose, CA*, Oct. 2004.
- [3] M. J. Neely and E. Modiano. Improving delay in ad-hoc mobile networks via redundant packet transfers. *Proc. of the Conference on Information Sciences and Systems*, Johns Hopkins University: March 2003.
- [4] R.L. Cruz and A. V. Santhanam. Hierarchical link scheduling and power control in multihop wireless networks. *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, Oct. 2002.
- [5] M. J. Neely, E. Modiano, and C. E. Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Journal on Selected Areas in Communications*, January 2005.
- [6] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. *Proceedings of IEEE INFOCOM*, 2001.
- [7] P. Gupta and P.R. Kumar. Critical power for asymptotic connectivity in wireless networks. *IEEE Conference on Decision and Control*, 1998.
- [8] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, vol. 46, no. 2:pp. 388–404, March 2000.
- [9] N. Bansal and Z. Liu. Capacity, delay and mobility in wireless ad-hoc networks. *IEEE Proceedings of INFOCOM*, April 2003.
- [10] M. Grossglauser and M. Vetterli. Locating nodes with ease: Last encounter routing in ad hoc networks through mobility diffusion. *IEEE Proceedings of INFOCOM*, April 2003.
- [11] E. Peravalov and R. Blum. Delay limited capacity of ad hoc networks: Asymptotically optimal transmission and relaying strategy. *IEEE Proceedings of INFOCOM*, April 2003.
- [12] A. Tsirigos and Z. J. Haas. Multipath routing in the presence of frequent topological changes. *IEEE Communications Magazine*, Nov. 2001.
- [13] E. Cohen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. *ACM Proceedings of SIGCOMM*, August 2002.
- [14] W. H. Yuen, R. D. Yates, and S-C Mau. Exploiting data diversity and multiuser diversity in noncooperative mobile infostation networks. *IEEE Proceedings of INFOCOM*, April 2003.
- [15] L. Kleinrock and J. A. Silvester. Optimum transmission radii in packet radio networks or why six is a magic number. *Proceedings of the National Telecommunications Conference*, Dec. 1978.
- [16] S. Toumpis and A. J. Goldsmith. Large wireless networks under fading, mobility, and delay constraints. *IEEE Proceedings of INFOCOM*, 2004.
- [17] A. El Gammal, J. Mammen, B. Prabhakar, and D. Shah. Throughput-delay trade-off in wireless networks. *IEEE Proc. of INFOCOM*, 2004.
- [18] X. Lin and N. B. Shroff. The fundamental capacity-delay tradeoff in large mobile ad hoc networks. *Purdue University Tech. Report*, 2004.
- [19] S. Ross. *Stochastic Processes*. John Wiley & Sons, Inc., New York, 1996.
- [20] F. P. Kelly. *Reversibility and Stochastic Networks*. Wiley, Chichester, 1984.
- [21] S. Ross. *Introduction to Probability Models*. Academic Press, 8th edition, Dec. 2002.
- [22] R. Gallager. *Discrete Stochastic Processes*. Kluwer Academic Publishers, Boston, 1996.
- [23] H. Daduna. *Queueing Networks with Discrete Time Scale*. Springer, 2001.
- [24] C.E. Koksal. *Providing Quality of Service over Electronic and Optical Switches*. PhD thesis, Massachusetts Institute of Technology, Laboratory for Information and Decision Systems (LIDS), Sept. 2002.



Michael J. Neely received B.S. degrees in both Electrical Engineering and Mathematics from the University of Maryland, College Park, in 1997. He was then awarded a 3 year Department of Defense NDSEG Fellowship for graduate study at the Massachusetts Institute of Technology, where he received an M.S. degree in EECS in 1999 and a Ph.D. in 2003. During the Summer of 2002, he worked in the Distributed Sensor Networks group at Draper Labs in Cambridge. In 2004 he joined the faculty of the Electrical Engineering Department at the University of Southern California, where he is currently an Assistant Professor. His research interests are in the areas of satellite and wireless networks, ad-hoc wireless networks, and queueing theory. Michael is a member of Tau Beta Pi and Phi Beta Kappa.



Eytan Modiano received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and PhD degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992-1993 while he was conducting research on security and performance issues in distributed network protocols.

Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he designed communication protocols for satellite, wireless, and optical networks and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. He joined the MIT faculty in 1999, where he is presently an Associate Professor in the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks.