Communication over a Wireless Network with Random Connections *

Radhika Gowaikar
Caltech
Pasadena, CA 91125

Bert Hochwald
Bell Laboratories
Murray Hill, NJ 07974

Babak Hassibi
Caltech
Pasadena, CA 91125

October 11, 2005

Abstract

We analyze a network of nodes in which pairs communicate over a shared wireless medium. We are interested in the maximum total aggregate traffic flow possible as given by the number of users multiplied by their data rate. Our model differs substantially from the many existing approaches in that the channel connections in our network are entirely random: we assume that, rather than being governed by geometry and a decay-versus-distance law, the strengths of the connections between nodes are drawn independently from a common distribution. Such a model is appropriate for environments where the first order effect that governs the signal strength at a receiving node is a random event (such as the existence of an obstacle), rather than the distance from the transmitter.

We show that the aggregate traffic flow as a function of the number of nodes n is a strong function of the channel distribution. In particular, for certain distributions the aggregate traffic flow is at least $\frac{n}{(\log n)^d}$ for some $d > 0$, which is significantly larger than the $O(\sqrt{n})$ results obtained for many geometric models. Our results provide guidelines for the connectivity that is needed for large aggregate traffic. We show how our model and distance-based models can be related in some cases.

1 Introduction

An early study of traffic flow in shared-medium wireless networks appears in the seminal work of Gupta and Kumar [11]. They show that in a grid network of n nodes on the plane having a deterministic power

*This work is supported in part by the National Science Foundation under grant nos. CCR-0133818 and CCR-0326554, by the David and Lucille Packard Foundation, and by Caltech's Lee Center for Advanced Networking.

scaling law, $O(\sqrt{n})$ transmitters can talk simultaneously to randomly-chosen receivers. Similar results for networks with randomly-placed nodes can also be obtained (see, for example, [10] for a recent account). Different models can yield somewhat different conclusions [1, 3, 5, 9, 12, 14, 15, 16, 17]; nevertheless, if we do not permit the transmitter/receiver pairs to approach one another [6], the model of a power decay law (as a function of distance) seems to yield a network in which the number of nodes that can talk simultaneously grows much slower than n. We wish to study networks with a different connectivity model.

The $O(\sqrt{n})$ result in [11] has the following heuristic explanation. If a node wishes to transmit directly to a randomly-chosen node (whose distance is approximately $O(\sqrt{n})$ away on average), it has two choices: talk directly, or talk through a series of hops. If it tries to talk directly, the transmitter generates energy in a circle of radius $O(\sqrt{n})$ around itself. However, this energy, which is seen by the intended receiver becomes interference for the $O(n)$ other nodes in the circle. Thus, some fraction of the entire network of n nodes is bathed in interference; an undesirable consequence. If it decides instead to talk through hops, the transmitting node can pass its message to a neighbor, who in turn passes it to a neighbor and so on for $O(\sqrt{n})$ hops to the intended receiver. This strategy limits interference to immediate neighbors but ties up $O(\sqrt{n})$ nodes in the hopping process. Although this turns out to be the best strategy, only $O(\sqrt{n})$ simultaneous messages can be passed before all n nodes in the network are involved.

We change the model of the wireless medium from a model based on distance to one based on randomness. In multi-antenna links, a linear increase in capacity (in the minimum of the number of transmit/receive antennas) is obtained when the channel coefficients between the transmit and receive antennas are independent Rayleigh-distributed random variables [4, 13]. It is therefore now generally believed that a rich scattering environment, once thought to be detrimental to point-to-point wireless communications, may actually be beneficial. We show that a similar effect may hold for the expected aggregate data traffic in a wireless network; certain forms of randomness can be helpful.

There are several reasons why one may choose a random model over one that is based on distance. While distance effects on signal strength are important for nodes that are very near or far from each other, many networks are designed with minimum and maximum distances in mind. Decay laws of the form $1/r^m$ for a fixed $m > 0$ may not be relevant for networks of small physical size. Additionally, through the use of automatic gain control, a radio often artificially mitigates distance effects unless the node is saturated (too close) or “dropped out” (too far). Many first-order signal-strength effects in such networks are then due to random fluctuations in the medium, such as Rayleigh and shadow fading. A distance-power model cannot readily account for shadow fading since signal strength at the receiver is determined more by the presence of an obstacle blocking the path to the transmitter than by distance.

We adopt the premise that randomness can have a first-order effect on the behavior of a network. We assume that the channels between nodes are drawn independently from an identical distribution. We allow the distribution of the channel between nodes to be arbitrary and allow it to vary with the number of nodes n. Our model covers environments where the the signal strength at a receiving node is governed primarily by a random event (such as the existence of an obstacle). We believe that the study of such wireless networks with random connections is important for three reasons: first, many real wireless networks have a substantial and dominant random component; second, we show that such networks may have qualitatively different traffic scaling laws than the scaling obtained in geometric models; finally, our results give insight into the connectivity that a network should have to allow large aggregate traffic flows.

In general, any realistic model of a large network should have a model of connectivity that has a balance of randomness and distance-based effects. For example, [8] uses a "radio model" to show that in the presence of obstructions and irregularities, channels become approximately uncorrelated with one another, and the probability of good links between nodes that are far apart increases in wireless local area networks (WLANs). The radio model in [8] essentially uses the same independence assumption that we do, but uses distance to determine the probability of a connection link. We show in Section 8.1.1 how to apply our traffic-flow conclusions to this radio model to determine a favorable distance between nodes.

1.1 Approach

We suppose that the connection strengths between the n nodes of the wireless network are drawn independently and identically from a given arbitrary distribution. In geometric networks such as [11] a node may communicate its message in hops to nearby neighbors so that it ultimately reaches the intended destination. In our random model, although there is no geometric notion of a near neighbor, we can find an equivalent of a near neighbor by introducing the notion of "good paths", where connections stronger than a chosen threshold β are called *good*. Transmissions to relays and destinations occur along only good paths. By figuratively drawing a graph whose vertices are all the nodes in the network, yet whose edges are only the good paths, we obtain a specific random graph model called $\mathcal{G}(n, p)$, where an edge between any pair of the n nodes exists with probability p. (In our case, p is simply the probability that the connection strength exceeds β_n.) $\mathcal{G}(n, p)$ is a very well-studied object and we leverage some of its known properties to establish disjoint routes between sources and their intended destinations. However since we are analyzing a wireless network, we must also account for the effects of interference between all nodes, including those that do not have good connections between them. Fortunately, our use of the goodness threshold β also makes the analysis of message-failures (due to interference and/or noise) tractable. Our analysis yields an achievable aggregate throughput which is a function of the chosen threshold β. A judicious choice of β can maximize this achievable throughput. To complement our achievability results, we also present on some upper bounds on aggregate throughput that show that our results are sometimes tight.

2 Model of Transmitted and Received Signals

We assume that the wireless network has narrowband flat-fading connections whose powers are independent and identically distributed (i.i.d.) according to an arbitrary distribution $f(\cdot)$. Thus, if $h_{i,j}$ is the connection between nodes i and j, then the $\gamma_{i,j} = |h_{i,j}|^2$ are i.i.d. random variables with marginal distribution $f(\gamma_{i,j})$. For maximum generality, we allow $f(\gamma) = f_n(\gamma)$ to be a function of the number of nodes n. As an example, consider

$$f(\gamma) = (1 - p) \cdot \delta(\gamma) + p \cdot \delta(\gamma - 1) \tag{1}$$

where $\delta(\cdot)$ is the Dirac delta-function. This distribution is a simple model of a shadow-fading environment where, for any pair of nodes, with probability p there exists a good connection between them (fading causes no loss), and with probability $1 - p$ there exists an obstruction (fading causes a complete loss). In a general network of n nodes, we may let $p = p_n$ be a function of n to represent changes in the geography or network topology as the network increases in size. Although $\gamma = 0$ and $\gamma = 1$ are the only possibilities in the distribution (1), we may also introduce values of γ that depend on n. Figure 1 pictorially displays an example of wireless terminals whose connections may obey the model (1).

The behavior of such a network varies dramatically with p. At the extreme of $p = 1$ no paths are ever blocked and all nodes are fully connected to each other. While this situation permits any node to readily talk to any other node in a single hop, the overall network throughput is low because talking pairs generate an enormous amount of interference for the remaining nodes. If many nodes try to talk simultaneously, the overall interference is overwhelming. At the other extreme of $p = 0$, everyone is in a deep fade; now interference is minimal. However, no nodes can talk at all (we assume a transmission power limit). Thus we have competing effects as a function of p: increasing p benefits the network by improving connectivity thus allowing for shorter hops, but hurts the network by increasing interference to other receivers. We are led to ask: what p is optimal? What is the resulting network aggregate traffic? Is this optimal p likely to be something we encounter naturally? If not, can we induce it artificially? We answer some of these questions

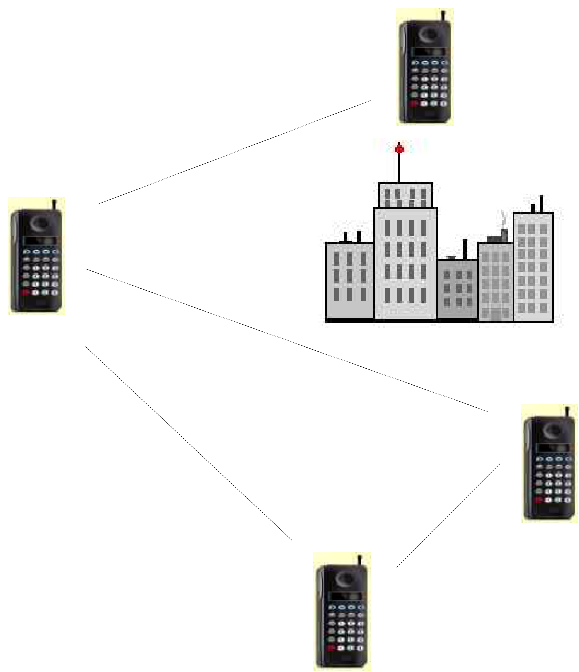

Figure 1: Nodes are able to establish connections with each other if there is no object in their path. Equation (1) models the presence of an object as a random event where each path has a connection of strength one with probability p, and otherwise has a connection of strength zero.

but, more generally, we look at how an arbitrary $f_n(\gamma)$ affects the traffic.

2.1 Detailed model

Let the network have n nodes labeled $1, \ldots, n$. Every pair of nodes $\{i, j\}$ ($i \neq j$) is connected by a channel that is denoted by the random variable $h_{i,j} = h_{j,i}$; there are $\binom{n}{2}$ channel random variables. The channel strengths, $\gamma_{i,j} = |h_{i,j}|^2$ are drawn i.i.d. according to the probability density function (pdf) $f_n(\gamma)$. Once drawn, these channel variables do not change with time.

Node i wishes to transmit signal x_i. We assume that x_i is a complex Gaussian random process with zero mean and unit variance. Each node is permitted a maximum power of P watts.

We incorporate interference and additive noise in our model as follows. Assume that k nodes $i_1, i_2, \ldots, i_k$ are simultaneously transmitting signals $x_{i_1}, x_{i_2}, \ldots, x_{i_k}$ respectively. Then, the signal received by node $j (\neq i_1, \ldots, i_k)$ is given by

$$y_j = \sum_{t=1}^{k} \sqrt{P} h_{i_t,j} x_{i_t} + w_j \tag{2}$$

where w_j represents additive noise. The additive noise variables $w_1, \ldots, w_n$ are i.i.d., drawn from a complex Gaussian distribution of zero mean and variance σ^2 ($w_j \sim \mathcal{CN}(0, \sigma^2)$). The noise is statistically independent of x_i.

2.2 Successful communication

In equation (2), suppose that only node i_1 wishes to communicate with node j and the signals $x_{i_2}, \ldots, x_{i_k}$ are interference. Then the signal-to-interference-plus-noise ratio (SINR) for node j is given by

$$\rho_j = \frac{P\gamma_{i_1,j}}{\sigma^2 + P\sum_{l=2}^{k} \gamma_{i_l,j}}$$

We assume that transmission is successful when the SINR exceeds some threshold ρ_0. If the SINR is less than ρ_0, we say that transmission is not possible.

3 Network Operation and Objective

We suppose that k nodes, denoted by $s_1, \ldots s_k$, are randomly chosen as sources. For every s_i, a destination node d_i is chosen at random, thus making k source-destination pairs. We assume that these $2k$ nodes are all distinct and therefore $k \leq n/2$. Source s_i wishes to transmit message M_i to destination d_i and has encoded it as signal x_i. We wish to see how many source-destination pairs may communicate simultaneously. The sources may talk directly to the destination nodes or may decide to communicate in hops through a series of relay nodes.

3.1 Communicating with hops

In general, we suppose that the source-destination pair (s_i, d_i) communicates using a sequence of relay nodes $r_{i,1}, r_{i,2}, \ldots, r_{i,h-1}$. ($h = 1, 2, \ldots$ represents the number of hops.) Define $r_{i,0} = s_i$ and $r_{i,h} = d_i$. The path from s_i to d_i is then $r_{i,0} = s_i, r_{i,1}, r_{i,2}, \ldots, r_{i,h-1}, r_{i,h} = d_i$. In time slot $t+1$ we have nodes $r_{1,t}, r_{2,t}, \ldots, r_{k,t}$ transmitting simultaneously to nodes $r_{1,t+1}, r_{2,t+1}, \ldots, r_{k,t+1}$ respectively. We ask that nodes $r_{1,t+1}, r_{2,t+1}, \ldots, r_{k,t+1}$ decode their respective signals $x_1, x_2, \ldots, x_k$ and transmit them to the next set of relay nodes in the $(t+2)$th time slot, and so on. A natural condition to impose is that the relay nodes that are receiving (or transmitting) messages in any time slot be distinct; the messages do not collide. In addition, we ask that relay nodes not receive and transmit at the same time. We refer to these conditions together as the property of *no collisions* in the rest of the paper. In general, we do not require $r_{i,t}$ to be distinct from $r_{i,t+1}$ for any i. This means that a relay can effectively hold on to a message in a time slot; hence h effectively represents the maximum number of hops needed for all the source-destination pairs.

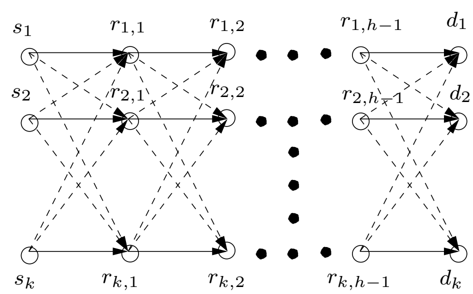

Figure 2: Schedule of relay nodes: Source s_i communicates with destination d_i using relays $r_{i,1}, \ldots, r_{i,h-1}$. The solid lines indicate intended transmissions and the dashed lines indicate potential interference. A schedule is valid if it meets the no-collision conditions that a node can receive or transmit at most one message in any time slot and that no node can transmit and receive simultaneously.

3.2 Throughput

With the above procedure, we have k simultaneous communications occurring in h time slots. Message M_i reaches the intended destination d_i successfully if it can be decoded by each relay $r_{i,t}$. Assume that a fraction $1-\epsilon$ of messages reach their intended destinations in this way. Then, we define the throughput as

$$T = (1-\epsilon)\frac{k}{h}\log(1+\rho_0), \tag{3}$$

where ρ_0 is the SINR threshold, and we are using the natural logarithm. Thus, $\log(1+\rho_0)$ is the sustainable throughput per user if the users do not collide. We multiply this factor by the number of non-colliding source-destination pairs k, divide by the number of hops, and subtract the fraction of dropped messages ϵ. The resulting throughput T depends on n and we sometimes add subscripts to the variables involved to indicate this: k_n, ϵ_n, $\rho_{0,n}$ and T_n. Typically, we force ϵ_n to go to zero as n grows. We demonstrate a scheme for choosing the relay nodes and analyze the throughput performance of this scheme. Thus, we give an achievability result for T_n. We now state this result.

4 Main Result

Theorem 1. *Consider a network on n nodes whose edge strengths are drawn i.i.d. from a probability distribution function $f_n(\gamma)$. Let $F_n(\gamma)$ denote the cumulative distribution function corresponding to $f_n(\gamma)$ and define $Q_n(\gamma) = 1 - F_n(\gamma)$. Choose any β_n such that $Q_n(\beta_n) = \frac{\log n + \omega_n}{n}$, where $\omega_n \to \infty$ as $n \to \infty$. Then there exists a positive constant α such that a throughput of*

$$T = (1-\epsilon_n)\, \alpha k_n(\beta_n)\frac{\log(nQ_n(\beta_n))}{\log n}\, \log\left(1 + \frac{a_n\beta_n}{\frac{\sigma^2}{P} + (k_n(\beta_n)-1)\mu_\gamma}\right) \tag{4}$$

is achievable for any positive a_n such that $a_n \leq 1$ and any $k_n(\beta_n)$ that satisfy the conditions:

1.

$$k_n(\beta_n) \leq \alpha n \frac{\log(nQ_n(\beta_n))}{\log n} \tag{5}$$

2.

$$\epsilon_n \leq \frac{a_n^2}{\alpha(1-a_n)^2}\frac{(k_n(\beta_n)-1)\sigma_\gamma^2}{(\frac{\sigma^2}{P} + (k_n(\beta_n)-1)\mu_\gamma)^2}\frac{\log n}{\log(nQ_n(\beta_n))} \to 0 \tag{6}$$

where μ_γ and σ_γ^2 are the mean and variance of γ respectively. The SINR threshold ρ_0 is given by $\frac{a_n\beta_n}{\frac{\sigma^2}{P}+(k_n(\beta_n)-1)\mu_\gamma}$.

The parameter β_n satisfying $Q_n(\beta_n) = \frac{\log n + \omega_n}{n}$ is the goodness threshold mentioned in Section 1.1. By figuratively drawing an edge when $\gamma > \beta_n$, we obtain a random graph that fits the well-studied model $\mathcal{G}(n,p)$. Condition (5) is needed to obtain a non-colliding schedule in this random graph. This issue is discussed in detail in Section 5. Once the schedule is obtained, we incorporate the effects of interference between non-colliding transmissions and provide an error analysis in Section 6. Condition (6) forces ϵ_n to go to zero. In Section 7 we combine the results of Sections 5 and 6 to prove the theorem. Note that the theorem indicates an achievable throughput and does not preclude that higher throughputs are possible.

Although not evident from the theorem statement at this time, it turns out that the optimum number of hops h grows at most logarithmically with n. The throughput therefore depends most strongly on the number of simultaneous transmissions k_n and the SINR threshold ρ_0.

The throughput expression (4) is general and accommodates an arbitrary $f_n(\gamma)$. The parameter k_n is the number of non-colliding simultaneous transmissions. We discuss the constant α and the parameter a_n later. The joint selection of β_n, k_n, and a_n that maximizes the achievable throughput (4) is not easily expressed in closed-form as a function of the pdf $f_n(\gamma)$. In general, these parameters need to be determined on a case-by-case basis. We show how to find the necessary parameters in Section 8 where we give several examples.

Since (4) holds for any k_n satisfying (5), we may choose k_n as large as possible (achieving equality in (5)) and optimize only over a_n and β_n. In fact, when $\frac{\sigma^2}{P} - \mu_\gamma \geq 0$, it is possible to show that the optimum k_n is the maximum possible. We hence state a more specific achievability result.

Corollary 1. *In the network of Theorem 1, if* $\frac{\sigma^2}{P} - \mu_\gamma \geq 0$ *the throughput (4) is maximized by choosing* k_n *as large as possible.*

5 Scheduling Transmissions

With a view to meeting a minimum SINR of ρ_0 at every relay node at every hop, we impose the condition that each transmitting link be stronger than some threshold β_n. We require that $\gamma_{r_{i,t},r_{i,t+1}} \geq \beta_n$, where β_n is a design parameter. We denote links that satisfy $\gamma_{i,j} \geq \beta_n$ as *good*. We require the path from s_i to d_i to use only good links.

The threshold β_n is a parameter that we may choose as a compromise between quantity and quality of the connections. By making β_n large we increase the quality of the link. However, if we make it too large we risk not being able to form an uninterrupted path of good links from the source to the destination. In this section, we determine the relation between β_n and the lengths of source-destination paths.

Define $p_n = \mathrm{P}(\gamma \geq \beta_n)$ (for convenience, we drop the subscript n in the rest of this section). Using our wireless communication network, we define a graph on n vertices as follows: For (distinct) vertices i and j of the graph, draw an edge (i,j) if and only if $\gamma_{i,j} \geq \beta_n$ in the network. Call the resulting graph

$G(n,p)$. The graph $G(n,p)$ then becomes an instance of a model called $\mathcal{G}(n,p)$ on n vertices in which edges are chosen independently and with probability p [2]. This graph shows the possible paths from the various sources to the various destinations using only good links, but does not show the possible interference encountered if these paths are used simultaneously. We examine this interference in Section 6.

Graphs taken from the model $\mathcal{G}(n,p)$ have many known properties. For instance, the values of p for which the graph is connected is well-characterized. As p increases the probability that the graph is connected goes to one. If $p = \frac{\log n + c + o(1)}{n}$ (where $c > 0$ need not be a constant) the probability of the graph being connected is $e^{-e^{-c}}$ [2]. This implies that there is a phase transition in the graph around $p = \frac{\log n}{n}$. For p less than this the probability of connectivity goes to zero rapidly and for p greater than this it goes to one rapidly. Another property that is well-studied is the *diameter*. The diameter of a graph is defined as the maximum distance between any two vertices of the graph, where the distance between two vertices is the minimum number of edges one has to traverse to go from one to the other. Results in [2] and [18] tell us that for p in the range of connectivity the diameter behaves like $\frac{\log n}{\log np}$. (It is also known that the average distance between two nodes has the same behavior.) This tells us that a message can be transmitted from one node to another using at most $\frac{\log n}{\log np}$ hops. What it leaves unanswered is the question of how to establish k such transmissions simultaneously and on non-colliding paths. In order to answer this question we invoke a relatively recent result regarding vertex-disjoint paths.

5.1 Scheduling using vertex-disjoint paths in $G(n,p)$

Two paths that do not share a vertex are called vertex-disjoint. Note that any two paths that are vertex-disjoint satisfy our "no-collisions" property; however, the reverse statement is not true. Thus, the vertex-disjoint condition is stronger than our requirement of non-colliding paths. For a set of k (disjoint) pairs of vertices (s_i, d_i), the question of whether there exists a set of vertex-disjoint paths connecting them is addressed in [19]. Their result states that with high probability, for every (sufficiently random) set of k pairs (s_i, d_i) and k not greater than $\alpha_1 n \frac{\log np}{\log n}$, where α_1 is a constant, there exists a set of vertex-disjoint paths. This result is within a constant of the best one can hope to achieve since the average distance between nodes in $\mathcal{G}(n,p)$ is $\frac{\log n}{\log np}$, and thus we can certainly have no more than $n\frac{\log np}{\log n}$ vertex-disjoint paths. Also stated in [19] is an algorithm that finds k paths using various random walk and flow techniques. Here we reproduce their main result.

Theorem 2. *Suppose that $G = G(n,p)$ and $p \geq \frac{\log n + \omega_n}{n}$, where $\omega_n \to \infty$. Then there exist two positive* *constants α_1, α_2 such that, with probability approaching 1, there are vertex-disjoint paths connecting s_i to d_i for any set of pairs*

$$F = \{(s_i, d_i)|s_i, d_i \in \{1, \dots, n\}, i = 1, \dots, k\}$$

satisfying

1. *The pairs in F for $i = 1, \dots, k$ are disjoint;*
2. *The total number of pairs, $k = |F|$, is not greater than $\alpha_1 n \frac{\log np}{\log n}$.*
3. *For every vertex $v \in \{1, \dots, n\}$, no more than a α_2-fraction of its set of neighbors, $N(v)$, are prescribed endpoints, that is $|N(v) \cap (S \cup D)| \leq \alpha_2 |N(v)|$, where $S = \{s_i\}$ and $D = \{d_i\}$.*

Furthermore, these paths can be constructed by an explicit randomized algorithm in polynomial time.

In fact, the existence of the paths is proved by stating and analyzing a randomized algorithm that finds them. However, we use this theorem only as an existence result to demonstrate achievable throughputs. Some comments about their randomized algorithm can be found in Sections 6 and 10.1.

In our communication network, Condition 1 that (s_i, d_i) be disjoint pairs is already met. The second imposes a restriction on how large k can be. Since the k source-destination pairs are chosen at random, the third condition is also met. (In fact, the third condition is imposed in [19] to prevent someone from choosing the (s_i, d_i) pairs in a particularly adversarial manner using knowledge of the graph structure.)

We can restate the theorem for our purposes.

Theorem 3. *Suppose that $G = G(n,p)$ and $p \geq \frac{\log n + \omega_n}{n}$, where $\omega_n \to \infty$. Then there exists a constant $\alpha > 0$ such that, with probability approaching 1, there are vertex-disjoint paths connecting s_i to d_i for any set of disjoint, randomly chosen source-destination pairs*

$$F = \{(s_i, d_i)|s_i, d_i \in \{1, \dots, n\}, i = 1, \dots, k\}$$

provided $k = |F|$ is not greater than $\alpha n \frac{\log np}{\log n}$.

The constant α in this theorem is the same α required in Theorem 1. It is not explicitly specified. We examine the lengths that these k paths can have in the following lemma.

Lemma 1. *Almost all of the $k = \alpha n \frac{\log np}{\log n}$ vertex-disjoint paths obtainable under Theorem 3 have lengths that grow no faster than $\frac{\log n}{\alpha \log np}$.*

Proof. Suppose that some fraction of paths, say $c_n k$ where $c_n > 0$ have average lengths of the form $\frac{\log n}{\log np}(1+\omega'_n)$ where ω'_n goes to infinity. Since there are n nodes in the network, we have

$$n \geq c_n k \times \frac{\log n}{\log np}(1+\omega'_n) = c_n \alpha n \frac{\log np}{\log n} \times \frac{\log n}{\log np}(1+\omega'_n) = c_n \alpha n (1+\omega'_n).$$

This implies that $1 \geq \alpha c_n (1+\omega'_n)$ and therefore c_n must go to zero. Therefore we conclude that at most a vanishing fraction of the k paths can have lengths that grow faster than $\frac{\log n}{\log np}$ and, asymptotically, all the paths have lengths that grow no faster than $\frac{\log n}{\alpha \log np}$. $\square$

Hence the number of hops h is (asymptotically) at most $\frac{\log n}{\alpha \log np}$. We use this fact in the error analysis in the following section.

6 Probability of Error

Consider a schedule of $k \leq \alpha n \frac{\log np}{\log n}$ non-colliding paths. Theorem 3 shows that such a schedule exists. One possible (but often impractical) way to obtain such a schedule is to use an exhaustive search that first lists all the paths between every source-destination pair and then randomly chooses a set that satisfies the vertex-disjoint property. Because we thereby choose a path based on vertices rather than edges, we are assured that any edges that might exist between vertices along one path to vertices along another are i.i.d. Bernoulli distributed with parameter p. We also conclude that the channel connections between nodes along different paths in the network are i.i.d. with distribution $f_n(\gamma)$.

More generally, randomized algorithms that choose non-colliding paths without using edge information between such paths also have the property of generating i.i.d. interference between the paths. An example of such a randomized algorithm that avoids an exhaustive search is [19].

We now consider the probability that a particular message fails to reach its intended destination. Destination d_i fails to receive message M_i if the SINR falls below ρ_0 at any of the h relay nodes $r_{i,1}, \ldots, r_{i,h} = d_i$. Denote by E_t the event that relay node $r_{i,t}$ does have an SINR greater than ρ_0. Note that the events $E_1, \ldots, E_h$ are identical. Therefore we have,

$$\mathrm{P}(M_i \text{ is received successfully}) = \mathrm{P}(\bigcap_{t=1}^{h} E_t) = 1 - \mathrm{P}(\bigcup_{t=1}^{h} \sim E_t) \geq 1 - \sum_{t=1}^{h} \mathrm{P}(\sim E_t) = 1 - h\mathrm{P}(\sim E_1) \quad (7)$$

where the inequality comes from the union bound. We now compute $\mathrm{P}(\sim E_1)$. This is the event that node

$r_{i,1}$ has an SINR lower than ρ_0

$$
\begin{aligned}
\mathrm{P}(\sim E_1) &= \mathrm{P}(\rho_{r_{i,1}} \leq \rho_0) \\
&= \mathrm{P}\left(\frac{P\gamma_{s_i,r_{i,1}}}{\sigma^2 + P\sum_{j\neq i}\gamma_{s_j,r_{i,1}}} \leq \rho_0\right) \\
&= \mathrm{P}\left(\sum_{j\neq i}\gamma_{s_j,r_{i,1}} \geq \frac{P\gamma_{s_i,r_{i,1}} - \rho_0\sigma^2}{P\rho_0}\right) \\
&\leq \mathrm{P}\left(\sum_{j\neq i}\gamma_{s_j,r_{i,1}} \geq \frac{P\beta_n - \rho_0\sigma^2}{P\rho_0}\right) \\
&= \mathrm{P}\left(\frac{1}{k-1}\sum_{j\neq i}\gamma_{s_j,r_{i,1}} - \mu_\gamma \geq \frac{P\beta_n - \rho_0\sigma^2}{(k-1)P\rho_0} - \mu_\gamma\right) \\
&\leq \mathrm{P}\left(\left|\frac{1}{k-1}\sum_{j\neq i}\gamma_{s_j,r_{i,1}} - \mu_\gamma\right| \geq \frac{P\beta_n - \rho_0\sigma^2}{(k-1)P\rho_0} - \mu_\gamma\right) \\
&\leq \frac{\sigma_\gamma^2/(k-1)}{(\frac{P\beta_n-\rho_0\sigma^2}{(k-1)P\rho_0} - \mu_\gamma)^2} \qquad (8)
\end{aligned}
$$

where the first inequality is because $\gamma_{s_i,r_{i,1}} \geq \beta_n$ and (8) comes from the Chebyshev inequality and the fact that the variance of $\frac{1}{k-1}\sum_{j\neq i}\gamma_{s_j,r_{i,1}}$ is $\sigma_\gamma^2/(k-1)$. The second inequality requires the condition $\frac{P\beta_n-\rho_0\sigma^2}{(k-1)P\rho_0} - \mu_\gamma \geq 0$, or

$$\rho_0 \leq \frac{\beta_n}{\frac{\sigma^2}{P} + (k-1)\mu_\gamma}. \qquad (9)$$

This condition on ρ_0 is intuitively satisfying: if we assume that k is large, then we expect the interference term in the denominator of the SINR to be approximately $(k-1)\mu_\gamma$. This would imply that setting the threshold ρ_0 to less than $\frac{\beta_n}{\frac{\sigma^2}{P}+(k-1)\mu_\gamma}$ would be sufficient to ensure that most hops would exceed this threshold.

We define ϵ_n to be the probability that the SINR threshold is not exceeded along one or more of the hops. From (7), $\epsilon_n \leq h\mathrm{P}(\sim E_1)$. We force $h\mathrm{P}(\sim E_1)$ to go to zero. From Lemma 1, h is at most $\frac{\log n}{\alpha \log np}$ and we have

$$\epsilon_n \leq h\mathrm{P}(\sim E_1) \leq \frac{\log n}{\alpha \log np}\frac{\sigma_\gamma^2}{(k-1)(\frac{P\beta_n-\rho_0\sigma^2}{(k-1)P\rho_0} - \mu_\gamma)^2} \qquad (10)$$

and we require the right-hand side to go to zero.

We mention that the inequality (10) requires γ to have a variance that does not go to infinity. If γ has infinite variance an alternative inequality is obtained from the Markov bound instead of the Chebyshev. The result is

$$\mathrm{P}(\sim E_1) \leq (k-1)\mu_\gamma \cdot \frac{P\rho_0}{P\beta_n - \rho_0\sigma^2}.$$

If needed, Theorem 1 can be modified to incorporate this inequality but we omit this modification since the Chebyshev bound is generally tighter.

7 Proof of Theorem 1

We now combine the results of Section 5 on the maximum number of non-colliding paths and Section 6 on the probability of successful transmission along these paths. We need $p = \mathrm{P}(\gamma \geq \beta_n) = Q_n(\beta_n) = \frac{\log n + \omega_n}{n}$ in order to do scheduling. In addition, we need:

1. To have non-colliding paths (Theorem 3)

$$k \leq \alpha n \frac{\log np}{\log n}$$

2. To meet the SINR threshold (equation (10))

$$\epsilon_n \leq \frac{\log n}{\alpha \log np} \frac{\sigma_\gamma^2}{(k-1)(\frac{P\beta_n - \rho_0\sigma^2}{(k-1)P\rho_0} - \mu_\gamma)^2} \to 0$$

3. To apply the Chebyshev inequality (equation (9))

$$\rho_0 \leq \frac{\beta_n}{\frac{\sigma^2}{P} + (k-1)\mu_\gamma}$$

To satisfy the third condition above we set

$$\rho_0 = \frac{a_n\beta_n}{\frac{\sigma^2}{P} + (k-1)\mu_\gamma}$$

where $0 \leq a_n \leq 1$. Substituting for this in the second condition, we get

$$\epsilon_n \leq \frac{a_n^2}{\alpha(1-a_n)^2} \frac{(k_n(\beta_n)-1)\sigma_\gamma^2}{(\frac{\sigma^2}{P} + (k_n(\beta_n)-1)\mu_\gamma)^2} \frac{\log n}{\log(nQ_n(\beta_n))} \to 0.$$

This and the first condition above are the only conditions on k. For any k satisfying these two conditions we get an achievable throughput. This gives us Theorem 1.

The theorem gives an achievable throughput as a function of β_n, a_n and k_n but does not attempt to optimize these parameters. Because ϵ_n goes to zero and h is determined by β_n, to find the optimum k we need to maximize $k\log(1+\rho_0) = k\log(1+\frac{a_n\beta_n}{\frac{\sigma^2}{P}+(k-1)\mu_\gamma})$ over k. In the particular case when $\frac{\sigma^2}{P} - \mu_\gamma$ is positive, the expression is non-decreasing in k (the first derivative is non-negative). Hence satisfying (5) with equality is optimum. This proves Corollary 1.

8 Examples and Applications

In this section we apply Theorem 1 to some particular channel distributions. Since, as in geometric models, the throughput is often interference-limited, we find that densities that lead to significant interference per transmitter generally underperform those that generate only a small amount of interference.

8.1 Shadow fading model

We revisit the model (1)

$$f_n(\gamma) = (1 - p_n)\delta(\gamma) + p_n\delta(\gamma - 1) \tag{11}$$

where $\delta(\cdot)$ is the Dirac delta-function. This pdf models the situation where strong shadow fading is present. The signal power is 0 in the presence of an obstruction and is 1 otherwise. We find the value of p that maximizes the throughput. (We drop the subscript n.) A natural choice for the goodness threshold β_n is 1, which gives $Q(\beta) = p$. We need to satisfy $p \geq (\log n + \omega_n)/n$ (where $\omega_n \to \infty$) in order to use Theorem 1.

Note that we have $\mu_\gamma = p$ and $\sigma_\gamma^2 = p(1 - p)$. Let us consider the case when $p = 1$. Then $\sigma_\gamma^2 = 0$ and equation (6) is always satisfied. The throughput expression becomes $T = \alpha k \log(1 + \frac{a_n}{\frac{\sigma^2}{P}+k-1})$ where $a_n \leq 1$. By considering different possible values of the optimizing k, is easy to check that the maximum throughput is no greater than a constant.

Let us consider the case when p is a constant other than 1. We consider three cases, $k = 1$, $k =$ constant $\neq 1$ and $k \to \infty$. In the first case (6) is satisfied easily but we get a constant throughput. In the second case, to satisfy (6), a_n must go to zero and the throughput also goes to zero. In the third case, the throughput becomes constant.

It remains to consider $p \to 0$. In this case, for sufficiently large n, the condition $\frac{\sigma^2}{P} - \mu_\gamma = \frac{\sigma^2}{P} - p \geq 0$ is satisfied. Therefore, according to Corollary 1 the maximum possible k achieves maximum throughput. Hence we consider $k = \alpha n \frac{\log np}{\log n}$. Since $p = \frac{\log n + \omega_n}{n}$, $k \to \infty$ and we may replace $k - 1$ by k in (6) and the SINR threshold. Since kp also goes to infinity, (6) becomes $\epsilon_n \leq \frac{a_n^2}{\alpha^2(1-a_n)^2} \frac{\log^2 n}{\log^2(np)} \frac{1}{n} \to 0$. Therefore a_n may be any positive constant $a < 1$. With this, the SINR threshold becomes $\rho_0 = \frac{a}{\frac{\sigma^2}{P} + \alpha np\frac{\log np}{\log n}} \approx \frac{a}{\alpha np\frac{\log np}{\log n}}$ which goes to zero. Thus $\log(1 + \rho_0) \approx \rho_0$ and we have $\frac{k}{h}\log(1 + \rho_0) = \frac{a\alpha}{p}\frac{\log np}{\log n}$. This is maximized when p is as small as possible, or $p = \frac{\log n + \omega_n}{n}$. The result is summarized in the Corollary.

Corollary 2. *Consider a network on n nodes where edge strengths are drawn i.i.d. from the distribution in*

(11). Then for large n the throughput is maximized for $p = \frac{\log n + \omega_n}{n}$ and is given by

$$T = \left(1 - \frac{a^2}{\alpha^2(1-a)^2}\frac{\log^2 n}{\log^2(\log n + \omega_n)}\frac{1}{n}\right) a\alpha \frac{\log(\log n + \omega_n)}{(\log n + \omega_n)\log n}\, n$$

as $n \to \infty$, where ω_n is any function going to infinity and $0 < a < 1$ and $\alpha < 1$ are constants.

This throughput is almost linear in n and requires the network to be sparsely connected; with a connection probability of $(\log n)/n$, each node is connected with only approximately $\log n$ other nodes. For example with $n = 1000$ nodes, we have $(\log n)/n = 0.0069$ and each node connects on average to only seven other nodes. Perhaps surprisingly, increasing or decreasing this connectivity has a detrimental effect. While it is clear that it is possible for a network to be under-connected, it is apparently also possible for a network to be over-connected. The simulations in Section 10.4 also demonstrate this effect.

8.1.1 Implications for a certain radio model

In [7, 8] a wireless connectivity model is introduced where the probability of a good link is expressed as

$$p(\hat{r}) = \frac{1}{2}\left[1 - \operatorname{erf}\left(3.07\frac{\log \hat{r}}{\xi}\right)\right] \tag{12}$$

where $\hat{r}$ is a (suitably normalized) distance between the transmitter and receiver and ξ is a parameter that depends on the degree of shadow fading and the distance pathloss exponent. Usually $\xi \in [0, 6]$ where large values indicate a strong shadow component. The links between different sources or destinations are modeled as statistically independent.

For nodes approximately $\hat{r}$ from each other, the model (12) is equivalent to our model of shadow fading (11) with $p = p(\hat{r})$. As we show in Section 8.1, maximum throughput is attained for $p \approx (\log n)/n$. The "equivalent distance" for nodes is found by solving

$$p = \frac{\log n}{n} = \frac{1}{2}\left[1 - \operatorname{erf}\left(3.07\frac{\log \hat{r}}{\xi}\right)\right]. \tag{13}$$

for $\hat{r}$. Nodes approximately this distance from each other then have the excellent throughput promised in Corollary 2. Because we cannot have a large network of nodes exactly equidistant from each other, equation (13) only has operational meaning if the link probability is relatively insensitive to the distance $\hat{r}$ when $p \approx (\log n)/n$. We show that it is.

As the number of nodes n increases, the optimum link-probability $(\log n)/n$ decreases, or, equivalently, the distance $\hat{r}$ between nodes increases. For large $\hat{r}$, we may approximate $\frac{1}{2}(1 - \operatorname{erf} x) \approx 1/(2\sqrt{\pi}x)\exp(-x^2)$, and thus (13) becomes

$$p = \frac{\log n}{n} = \frac{\xi}{10.88 \log \hat{r}} e^{-3.07 \log^2 \hat{r}/\xi}.$$

The sensitivity of p as a function of $\hat{r}$ is very low when p is small. We show this in Figure 3, where we display p versus $\hat{r}$ for various values of ξ. The dotted lines in the figure shows the approximate optimum operating point p for networks with 100 and 1000 nodes. We see that the optimum p is generally very small and relatively insensitive to $\hat{r}$, and the best network performance is generally therefore obtained when the nodes are relatively far apart from one another, with a wide range of acceptable distances. This suggests that a large high-throughput network of nodes with optimum (small) p is possible.

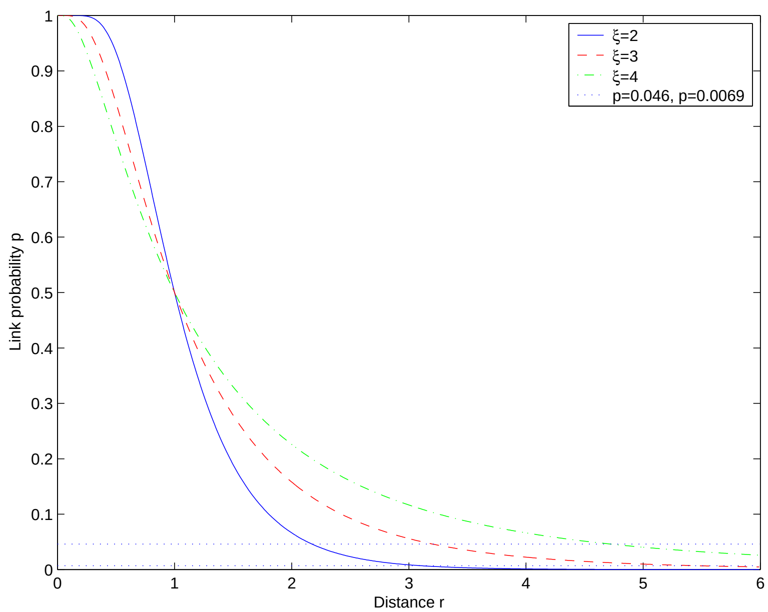

Figure 3: Link probability p versus distance $\hat{r}$ as given by (13) for $\xi = 2, 3, 4$. Also shown are dotted lines at $p = (\log 100)/100 \approx 0.046$ and $p = (\log 1000)/1000 \approx 0.0069$ indicating the optimum throughput point for shadow-fading with 100 and 1000 nodes respectively. As a function of $\hat{r}$, p is relatively insensitive for large $\hat{r}$.

We comment that the authors in [8] also consider how shadow fading can reduce the hop-count in a network and they use some graph-theoretic concepts in their arguments. They do not, however, attempt to obtain a throughput result by finding simultaneous non-colliding paths, nor do they incorporate the detrimental effects of interference to show that a network can be "too connected".

8.2 An exponential density

Let $f_n(\gamma) = e^{-\gamma}$. For this pdf, the mean and variance are constant, independent of n. We consider the three cases, $k = 1$, $k =$ constant $\neq 1$ and $k \to \infty$. In the first case, (6) is satisfied and a_n can be any positive constant $a < 1$; in the second case we need $a_n \to 0$ to satisfy (6). However, in either case, since we have $Q_n(\beta_n) = e^{-\beta_n} = \frac{\log n + \omega_n}{n}$, we have $\beta_n \leq \log n$. Also, simply using the fact that $Q_n(\beta_n) \leq 1$ for the optimum β_n, we get that the throughput $T \leq \alpha k \log(1 + \frac{a_n \log n}{\frac{\sigma^2}{P} + (k-1)})$.

We may compare the results of [22], where a multi-antenna broadcast channel with M transmit antennas to n users (each with N receive antennas) is studied. For M growing slower than $\log n$ (in particular, for constant M) the throughput scales as $M \log\log nN$. This double-logarithmic growth in n is very similar to the achievable throughput obtained above where k plays the role of the number of antennas M.

Let us consider the case when $k \to \infty$. We replace $k - 1$ by k in the throughput expression of (4), the condition of (6) and the SINR threshold. With this, (4) becomes an increasing function of k and hence the maximum permissible k becomes optimal. Therefore we have $k = \alpha n \frac{\log n e^{-\beta_n}}{\log n}$. This value of k also satisfies (6) with constant a_n. We obtain the optimum $\rho_0 = \frac{a\beta}{\frac{\sigma^2}{P} + \alpha n \frac{\log(ne^{-\beta})}{\log n}}$ and $\epsilon_n = \frac{a^2}{\alpha^2(1-a)^2} \frac{\log^2 n}{n \log^2(ne^{-\beta})}$ which goes to zero because $\beta \leq \log n$. This gives us a throughput of

$$T = (1 - \epsilon_n)\alpha^2 n \frac{\log^2(ne^{-\beta})}{\log^2 n} \log\left(1 + \frac{a\beta}{\frac{\sigma^2}{P} + \alpha n \frac{\log(ne^{-\beta})}{\log n}}\right).$$

Ignoring the ϵ_n term, we find the β that maximizes this expression. Since $\beta \leq \log n$, ρ_0 goes to zero and we can approximate $\log(1 + \rho_0) \approx \rho_0$. Hence, we have to maximize $\frac{\log^2(ne^{-\beta})}{\log^2 n} \frac{a\beta}{\frac{\sigma^2}{P} + \alpha n \frac{\log(ne^{-\beta})}{\log n}}$ over β. The term σ^2/P is negligible compared with the remaining term in the denominator (which goes to infinity) and therefore this expression becomes $\frac{\alpha \log(ne^{-\beta})}{\log n} \frac{a\beta}{n}$, and the maximizing β is $\log n/2$. We note that $p = \frac{1}{\sqrt{n}}$, the number of hops is 2, and ϵ_n becomes $\frac{a^2}{\alpha^2(1-a)^2} \frac{4}{n}$. The throughput becomes $T = \left(1 - \frac{a^2}{\alpha^2(1-a)^2} \frac{4}{n}\right) \frac{a\alpha \log n}{4}$. For large n this throughput is larger than the throughput $T = \alpha k \log(1 + \frac{a_n}{\frac{\sigma^2}{P} + k - 1})$ obtained for constant k. (For small n the latter throughput may sometimes be larger.)

Corollary 3. *Consider a network on n nodes where edge strengths are drawn i.i.d. from a distribution $f_n(\gamma) = e^{-\gamma}$. Then a throughput of*

$$T = \left(1 - \frac{a^2}{\alpha^2(1-a)^2} \frac{4}{n}\right) \frac{a\alpha \log n}{4}$$

is achievable as $n \to \infty$ where $\alpha < 1$, $a < 1$ are constants.

We see that a random network dominated by an exponential pdf has a throughput that scales only logarithmically with n. This network has good connectivity since the number of hops is small, but is also unfortunately dominated by interference. Thus, only few transmissions can occur simultaneously. We show in Section 9 that this throughput is tight to first-order in n.

8.3 Density obtained from a decay law

In this example we construct a pdf from the marginal density of the channel strengths in a geometric model. For every node, the channel coefficients to the remaining nodes follow a deterministic law based on distance. If we group these coefficients according to their magnitude γ, we obtain a certain number of coefficients whose magnitude falls in the interval $(\gamma, \gamma + d\gamma)$. We seek a probability density function whose average number of magnitudes matches this deterministic law.

In an actual geometric model the distribution of channel magnitudes depends on the location of the nodes. We make a simplifying assumption: We suppose that the nodes are in a circular disk and consider the node at the center of the disk to derive the density. We thereby ignore the effects of the disk boundary. We assume the nodes are dropped with density Δ (nodes per unit area) but ensuring a minimum distance of d from the center. The area of the entire disk is n/Δ.

In deriving the density of the channel coefficients, we use a power law of the form $g(r)$, where a node transmitting with power P is received by another node at distance r with power $Pg(r)$. We assume that $g(\cdot)$ is monotonically decreasing. The most significant difference between our model and the standard geometric model is in the independence of the channel coefficients in our model that does not exist in the geometric model. The geometric model has a correlation structure in the coefficients where channels of similar strength are clustered in rings around the center node. In our model, coefficients of similar strength, although the same in number as the geometric model, are distributed randomly and not necessarily geometrically colocated.

Consider a node at the center of the disk transmitting at power P. The fraction of nodes receiving power $\leq \gamma P$ is given by $1 - \frac{\Delta}{n} 2\pi((g^{-1}(\gamma))^2 - d^2)$ where $\gamma \in [g(\sqrt{\frac{n}{2\pi\Delta} + d^2}), g(d)]$ In particular, if we have a decay law of the form $g(r) = \frac{1}{r^m}$, this tells us that the fraction of nodes receiving power $\leq \gamma P$ is given by

$$1 - \frac{\Delta}{n} 2\pi(\frac{1}{\gamma^{2/m}} - d^2)$$

for $\gamma \in [\left(\frac{2\pi\Delta}{n+2\pi\Delta d^2}\right)^{m/2}, \frac{1}{d^m}]$.

This is a cumulative distribution function and by differentiating it with respect to γ we obtain the pdf for the edge strengths seen by the central node as

$$f_n(\gamma) = \frac{4\pi\Delta}{nm}\frac{1}{\gamma^{1+\frac{2}{m}}}, \qquad \gamma \in \left[\left(\frac{2\pi\Delta}{n+2\pi\Delta d^2}\right)^{m/2}, \frac{1}{d^m}\right], \qquad m > 0. \tag{14}$$

We assume that connections are drawn i.i.d. from this distribution.

We apply our results to this network. We have $Q_n(\beta_n) = \frac{\Delta}{n}2\pi(\frac{1}{\beta_n^{2/m}} - d^2)$. Since we need $Q_n(\beta_n) = p_n = \frac{\log n+\omega_n}{n}$, we have $\beta_n \leq (d^2 + \frac{\log n+\omega_n}{2\pi\Delta})^{-m/2} \approx (2\pi\Delta/(\log n + \omega_n))^{m/2}$. For different values of m, the mean and variance of γ can be evaluated. For large n, these are:

$$\mu_\gamma = \begin{cases} \frac{2(2\pi\Delta)^{m/2}}{2-m}\, n^{-m/2} & m < 2 \\ 2\pi\Delta\, \frac{\log n}{n} & m = 2 \\ \frac{4\pi\Delta}{(m-2)d^{m-2}}\, \frac{1}{n} & m > 2. \end{cases} \tag{15}$$

$$\sigma_\gamma^2 = \begin{cases} (2\pi\Delta)^m(\frac{1}{1-m} - \frac{4}{(2-m)^2})\, n^{-m} & m < 1 \\ 2\pi\Delta\, \frac{\log n}{n} & m = 1 \\ \frac{2\pi\Delta}{(m-1)d^{2(m-1)}}\, \frac{1}{n} & m > 1. \end{cases} \tag{16}$$

Since the mean goes to zero in each case, the condition $\frac{\sigma^2}{P} - \mu_\gamma \geq 0$ is met for sufficiently large n. Therefore we can use Corollary 1. It turns out that, in each case, the optimum β_n makes p_n as small as possible, which is $\beta_n = (2\pi\Delta/(\omega_n + \log n))^{m/2}$. Values of ϵ_n and the throughput can be calculated for each of these cases. We omit the details and simply state these results in the following corollary.

Corollary 4. *Consider a network on n nodes where edge strengths are drawn i.i.d. from the distribution*

$$f_n(\gamma) = \frac{4\pi\Delta}{nm}\frac{1}{\gamma^{1+\frac{2}{m}}}, \qquad \gamma \in \left[\left(\frac{2\pi\Delta}{n+2\pi\Delta d^2}\right)^{m/2}, \frac{1}{d^m}\right], \qquad m > 0$$

Then the following values of ϵ_n and throughputs are achievable:

$$\epsilon_n \leq \begin{cases} \frac{a^2}{\alpha^2(1-a)^2}(\frac{(2-m)^2}{4(1-m)} - 1)\frac{\log^2 n}{\log^2(\log n+\omega_n)}\frac{1}{n} & m < 1 \\ \frac{a^2}{4(1-a)^2}\frac{\log^3 n}{\log^2(\log n+\omega_n)}\frac{1}{\alpha^2 n} & m = 1 \\ \frac{a^2(2\pi\Delta)^{1-m}(2-m)^2}{4(1-a)^2(m-1)d^{2(m-1)}}\frac{\log^2 n}{\log^2(\log n+\omega_n)}\frac{1}{\alpha^2 n^{2-m}} & 1 < m < 2 \\ \frac{a^2}{2\pi\Delta(1-a)^2 d^2}\frac{1}{\alpha^2\log^2(\log n+\omega_n)} & m = 2 \\ \frac{1}{w_n^2}\frac{2\pi\Delta P^2}{(m-1)d^{2(m-1)}\alpha\sigma^4} & m > 2. \end{cases} \tag{17}$$

$$T = \begin{cases} (1-\epsilon_n)\frac{a(2-m)\alpha}{2}\frac{\log(\log n+\omega_n)}{\log n(\log n+\omega_n)^{m/2}}\, n^{m/2} & m<1 \\ (1-\epsilon_n)\frac{a\alpha}{2}\frac{\log(\log n+\omega_n)}{\log n(\log n+\omega_n)^{1/2}}\, n^{1/2} & m=1 \\ (1-\epsilon_n)\frac{a(2-m)\alpha}{2}\frac{\log(\log n+\omega_n)}{\log n(\log n+\omega_n)^{m/2}}\, n^{m/2} & 1<m<2 \\ (1-\epsilon_n)a\alpha\frac{\log(\log n+\omega_n)}{\log^2 n(\log n+\omega_n)}\, n & m=2 \\ (1-\epsilon_n)\frac{P\alpha^2(2\pi\Delta)^{m/2}}{\sigma^2 w_n}\frac{\log^2(\log n+\omega_n)}{\log^2 n(\log n+\omega_n)^{m/2}}\, n & m>2. \end{cases} \tag{18}$$

where $a < 1$ and $\alpha < 1$ are constants and ω_n and w_n are functions going to infinity.

We see that almost linear throughput can be obtained for $m \geq 2$. This differs substantially from the $O(\sqrt{n})$ results obtained for the structured deterministic model with the same decay law. Our results show that it is not the marginal distribution of the power that impedes the throughput in a geometric power-decay network, but rather the spatial distribution of these powers.

8.4 A heavy tail distribution

Consider a network on n nodes where edge strengths are drawn i.i.d. from $f_n(\gamma) = \frac{c}{1+\gamma^4}, \gamma \geq 0$ where c is such that $f_n(\gamma)$ integrates to 1. Clearly, the mean and variance of this distribution are constant with respect to n. We consider the three cases $k = 1$, $k =$ constant $\neq 1$ and $k \to \infty$. In the first case (6) is satisfied easily and we can set a_n to any positive constant $a < 1$; in the second we need $a_n \to 0$. In either case, the optimum β behaves like $n^{1/6}$ which gives $p = \frac{3c}{n^{1/2}}$ and a throughput of $\frac{\alpha}{2}k\log(1+\frac{a_n n^{1/6}}{\frac{\sigma^2}{P}+(k-1)\mu_\gamma})$. For $k \to \infty$, we replace $k-1$ with k in (6), (5) and (4). With this, one can determine that the optimum β_n maximizes $\frac{\beta_n}{h}$ or $\beta_n \log(nQ_n(\beta_n))$ while still satisfying $Q_n(\beta_n) = p_n = \frac{\log n+\omega_n}{n}$. The smallest value of p_n turns out to be optimal, leading to $\beta_n = \frac{n^{1/3}}{(\log n+\omega_n)^{1/3}}\frac{c^{1/3}}{3^{1/3}}$. We have the following corollary.

Corollary 5. *Consider a network on n nodes where edge strengths are drawn i.i.d. from the distribution $f_n(\gamma) = \frac{c}{1+\gamma^4}, \gamma \geq 0$. The throughput is then*

$$\begin{aligned} T &= \left(1 - \frac{a^2\sigma_\gamma^2}{\alpha^2\mu_\gamma^2(1-a)^2}\frac{\log^2 n}{\log^2(\log n+\omega_n)}\frac{1}{n}\right)\frac{a(c/3)^{1/3}\alpha}{\mu_\gamma}\frac{\log(\log n+\omega_n)}{\log n(\log n+\omega_n)^{1/3}}\, n^{1/3} \\ &\approx \frac{a(c/3)^{1/3}\alpha}{\mu_\gamma}\frac{\log\log n}{\log^{4/3} n}\, n^{1/3}. \end{aligned}$$

8.5 A distribution with constant mean and variance

The throughput of the previous example as well as that of the exponential density of Section 8.2 can be derived from the general case where $f_n(\gamma)$ has constant mean and variance as a function of n. Consider the cases $k = 1$, $k = \text{constant} \neq 1$ and $k \to \infty$. When $k = 1$ (6) is satisfied and we may choose a_n to be any positive constant $a < 1$. In the second case, to satisfy (6) we need $a_n \to 0$. In either case, the optimum β_n is found by maximizing $T = \frac{\log nQ_n(\beta_n)}{\log n} \log(1 + \frac{a_n P \beta_n}{\sigma^2})$ subject to the condition $p_n = Q_n(\beta_n) \geq \frac{\log n + \omega_n}{n}$ where $a_n \leq 1$.

In the third case $k \to \infty$ condition (6) is always satisfied for constant a_n and the optimum $\rho_0 = \frac{a\beta_n}{\frac{\sigma^2}{P} + k\mu_\gamma}$ goes to zero. We have $\log(1 + \rho_0) \approx \rho_0 \approx \frac{a\beta_n}{k\mu_\gamma}$. Now, $T = \frac{k}{h} \log(1 + \rho_0) \approx \frac{a}{\mu_\gamma} \frac{\beta_n}{h} = \frac{a\alpha}{\mu_\gamma} \frac{\beta_n \log(nQ_n(\beta_n))}{\log n}$. Therefore, to maximize the throughput, we need to maximize $\frac{\log(nQ_n(\beta_n))}{\log n} \frac{a\alpha\beta_n}{\mu_\gamma}$ subject to the condition that $p_n = Q_n(\beta_n) \geq \frac{\log n + \omega_n}{n}$.

Comparing the two objective functions using $\log(1 + x) \leq x$, we may always choose the case where $k \to \infty$ to get the larger throughput (for large n). This gives us $\epsilon_n \leq \frac{a^2\sigma_\gamma^2}{\alpha^2(1-a)^2\mu_\gamma^2} \frac{\log^2 n}{\log^2(nQ_n(\beta_n))} \frac{1}{n}$. (Here $k - 1$ has been replaced by k.) This result mirrors the derivations used in the exponential and heavy tail distributions. We comment that although we obtain the best throughput by letting $k \to \infty$ as $n \to \infty$, for small n a small k may give a numerically higher throughput than a large one.

Corollary 6. *Consider a network on n nodes where edge strengths are drawn i.i.d. from a distribution $f_n(\gamma)$ where the mean μ_γ and variance σ_γ^2 of γ are independent of n. Then the throughput is given by*

$$T = \left(1 - \frac{a^2\sigma_\gamma^2}{\alpha^2(1-a)^2\mu_\gamma^2} \frac{\log^2 n}{\log^2(nQ_n(\beta_n))} \frac{1}{n}\right) \frac{a\alpha}{\mu_\gamma} \frac{\beta_n \log(nQ_n(\beta_n))}{\log n}$$

and the optimum β_n maximizes $\beta_n \log(nQ_n(\beta_n))$ while satisfying $Q_n(\beta_n) \geq \frac{\log n + \omega_n}{n}$.

Perhaps surprisingly, distributions with constant mean and variance, while allowing us to apply Corollary 6, can have widely different throughputs. For example, both the exponential and heavy tail distributions examined earlier have constant mean and variance but the throughput in the exponential is logarithmic in n while the throughput in the heavy-tail is roughly $n^{1/3}$.

8.6 Tradeoff between k and ρ_0

In most of the examples above we notice that the optimum k goes to infinity; hence the optimum $\rho_0 = \frac{a\beta_n}{\frac{\sigma^2}{P} + (k-1)\mu_\gamma}$ goes to zero. In these cases we approximate $\log(1 + \rho_0)$ by ρ_0. In addition, if $k\mu_\gamma$ goes to infinity, we can further approximate ρ_0 as $\frac{a\beta_n}{k\mu_\gamma}$. In this case, we have $\frac{k}{h} \log(1 + \rho_0) \approx \frac{a\beta_n}{h\mu_\gamma}$. This expression depends only on β_n and is independent of k and ρ_0. We can therefore increase (decrease) k, thus decreasing (increasing) $\rho_0 = \frac{a\beta_n}{\frac{\sigma^2}{P} + (k-1)\mu_\gamma}$ and (as long as $k\mu_\gamma \to \infty$) the throughput remains unaffected. Hence it is sometimes possible to trade off the number of simultaneously communicating source-destination pairs with the SINRs at which they communicate without affecting the aggregate throughput.

9 Upperbounds

Our method of finding the throughput relies on finding good edges along which the desired communication can take place. This method does not preclude other methods from possibly doing better. In the cases where the throughput is of the form $\frac{n}{\log^d n}$ the optimal throughput cannot be better by more than the factor $\log^d n$ because the maximum throughput cannot scale more than linearly (unless the channel density is somehow chosen such that the maximum received power increases as the number of nodes increases – we exclude such densities here).

However, when the throughput we compute turns out to be of the order of n^d for $d < 1$, or $\log n$ as with the exponential density, it is not clear that we cannot do better. In this section we present an approach to computing an upperbound on throughput that shows that we sometimes cannot do better.

The throughput is given by $(1-\epsilon)\frac{k}{h}\log(1+\rho_0)$. We ignore the h in the denominator and find an upperbound for $k\log(1+\rho_0)$. Thus, we allow ourselves to choose k source-destination pairs from a given network and find the highest SINR threshold that can be met for all of them simultaneously. This is equivalent to finding a bound for the best single hop communication. Clearly, by doing this, our achievability results are certain to be at least a factor of h away from the upperbound. However, we know that h can be no larger than $\frac{\log n}{\log(\log n+\omega_n)}$, which is often a small factor.

There are $\binom{n}{k}\binom{n-k}{k}k!$ ways of choosing k source-destination pairs in a network. Assume that a threshold ρ_0 is fixed. Then, for a randomly drawn set of source-destination pairs, there is a probability, say p_s, that a received message satisfies the SINR threshold and is decoded successfully. The probability that all k received messages satisfy the threshold is p_s^k. Therefore, for a given pair (k,ρ_0), the expected number of sets of k source-destination pairs that satisfy the threshold ρ_0 is

$$M_n(k,\rho_0) = \binom{n}{k}\binom{n-k}{k}k!p_s^k.$$

Note that p_s depends on ρ_0, k and the pdf $f_n(\gamma)$ from which the connections are drawn. We say that a (k,ρ_0) pair is feasible if there exists at least one set of k source-destination pairs such that each of the k SINRs exceeds ρ_0. The probability that a particular (k,ρ_0) pair is feasible can be bounded as follows.

$$\begin{aligned}
\mathrm{P}((k,\rho_0) \text{ is feasible}) &= \mathrm{P}(\text{\# of } k\text{-pairs that satisfy the threshold } \rho_0 \text{ is} \geq 1)\\
&\leq \mathrm{E}(\text{\# } k\text{-pairs that satisfy the threshold } \rho_0)\\
&= M_n(k,\rho_0)
\end{aligned}$$

where the Markov inequality is used.

If for a particular choice of (k, ρ_0) we have $M_n(k, \rho_0)$ going to zero then that choice is infeasible. Otherwise (k, ρ_0) may be feasible. We can thereby characterize all (k, ρ_0) pairs that may be feasible. The largest value of $k\log(1+\rho_0)$ taken over these pairs gives us an upperbound on the throughput.

Note that this approach is general and can be used for any pdf, but requires a calculation of

$$p_s = \mathrm{P}\left(\frac{P\gamma_1}{\sigma^2 + P\sum_{i=2}^{k}\gamma_i} \geq \rho_0\right).$$

where all the channel coefficients in the SINR expression are drawn i.i.d. according to $f_n(\gamma)$. For certain densities, such as the exponential, we may compute p_s and get an upperbound as follows.

If $f_n(\gamma) = e^{-\gamma}$, then

$$p_s = \mathrm{P}\left(\frac{P\gamma_1}{\sigma^2 + P\sum_{i=2}^{k}\gamma_i} \geq \rho_0\right) = \frac{e^{-\frac{\sigma^2}{P}\rho_0}}{(1+\rho_0)^{k-1}}.$$

With this,

$$M_n(k,\rho_0) = \binom{n}{k}\binom{n-k}{k}k!\frac{e^{-\frac{\sigma^2}{P}k\rho_0}}{(1+\rho_0)^{k(k-1)}}.$$

We now want to characterize (k, ρ_0) pairs for which $M_n(k, \rho_0)$ does not go to zero. We have

$$M_n(k,\rho_0) = \frac{n!}{(n-2k)!k!}p_s^k \leq \frac{n!}{(n-2k)!}p_s^k \leq n^{2k}\frac{e^{-\frac{\sigma^2}{P}k\rho_0}}{(1+\rho_0)^{k(k-1)}} \leq \left(n^2\frac{1}{(1+\rho_0)^k}\right)^k = e^{k(2\log n - k\log(1+\rho_0))}.$$

If k goes to infinity (with n) and $2\log n - k\log(1+\rho_0)$ is negative then $M_n(k, \rho_0)$ goes to zero. Therefore, for k going to infinity, we have $k\log(1+\rho_0) \leq 2\log n$ as a bound on the throughput. If k is constant, it is easy to see that $1+\rho_0$ cannot grow faster than n^2, hence the throughput is again limited by $k\log n^2 = 2k\log n$ where k is now a constant. Thus we have shown an upperbound of $c\log n$ on the throughput. This happens to coincide (to within a constant) with the throughput obtained in our achievability result (Section 8.2). In our scheme it turns out that using two hops is optimal for any n. Hence, although the upperbound derived here is on $k\log(1+\rho_0)$, it matches the achievability result for $\frac{k}{h}\log(1+\rho_0)$ very closely.

10 Discussions, Simulations and Conclusions

Theorem 1 gives a very specific achievability result but equation (4) involves a constant α that is not explicit. This constant has its origins in Theorem 3 where the number of vertex-disjoint paths is computed. When we are confronted with a specific network with a finite number of nodes n, we would like an explicit estimate of the number of non-colliding paths. In this section we provide such an estimate; we also briefly introduce the notion of “bad” edges, discuss decentralized algorithms for attaining our achievability results, and provide computer simulations of some of the networks analyzed in Section 8.

10.1 Non-colliding paths

In section 5 we use a result of [19] to establish the existence of non-colliding paths. In this section, we present a constructive method of obtaining these paths and analyze the expected number of non-colliding paths thereby obtained. The algorithm we present is used extensively in Section 10.4.

We begin by choosing nodes $1, \ldots, n/2$ as source nodes and nodes $n/2+1, \ldots, n$ as their respective destination nodes. For the first source-destination pair, a shortest path connecting them (using only links that exceed β) is found. This is done using a standard breadth-first search algorithm [20] in which a rooted tree is constructed. All of the nodes begin by being “undiscovered”. The source node acts as the root of the tree (at depth zero) and is labeled as “discovered”. We then find all the nodes that are its neighbors and call them discovered. These are at distance one from the source and hence at depth one in the breadth-first search tree. The nodes at depth one are then processed successively. All of the neighbors of each node that are still undiscovered are put in the tree at depth two and their labels are changed to discovered. The process continues till there are no undiscovered nodes. Clearly, each node appears at most once in the tree. A shortest path from the source (root) to the destination is obtained by simply finding that node in the tree and moving up the tree to the source node. If the destination does not appear in the tree it has no path to the source.

Once the shortest path for the ith source-destination pair is established it is recorded and all n nodes are relabeled as “undiscovered”; the entire process is repeated to find the shortest path for the $(i+1)$st source-destination pair. This is done till paths are found for all $n/2$ pairs.

We then eliminate colliding paths on this list, starting with the first source-destination pair. If a node used on the path between s_1 and d_1 collides with a node on some other path, we eliminate path 1, otherwise we keep it. We proceed in order and eliminate the ith path if it collides with any of paths $i+1, i+2, \ldots, n/2$

and keep it otherwise. Note that since we start with shortest paths, a relay never appears more than once on a particular path.

Let us bound the probability that paths i and j collide for $i \neq j$. Without loss of generality we can set $i = 1$ and $j = 2$. We now have

$$\begin{aligned}
& \mathrm{P}(\text{path 1 collides with path 2}) \\
= \; & \mathrm{P}\left((s_1 = r_{2,1}) \cup \bigcup_{j=1}^{h-1} (r_{1,j} = r_{2,j-1} \cup r_{1,j} = r_{2,j} \cup r_{1,j} = r_{2,j+1}) \cup (d_1 = r_{2,h-1}) \right) \\
\leq \; & \mathrm{P}(s_1 = r_{2,1}) + \sum_{j=1}^{h-1} \mathrm{P}(r_{1,j} = r_{2,j-1}) + \sum_{j=1}^{h-1} \mathrm{P}(r_{1,j} = r_{2,j}) + \sum_{j=1}^{h-1} \mathrm{P}(r_{1,j} = r_{2,j+1}) + \mathrm{P}(d_1 = r_{2,h-1}) \\
= \; & \frac{3h-1}{n-2} \qquad (19)
\end{aligned}$$

The inequality is a standard union bound and the last equality is because the $h-1$ relay nodes on the ith path are drawn uniformly at random from from the set of all nodes of the graph (excluding s_i and d_i). (We assume that the algorithm that chooses the shortest path for (s_i, d_i) does not use any knowledge of the previously chosen $i-1$ paths.)

Denote by D_i the event of keeping the ith path. This event comprises the intersection of the events that the ith path does not collide with the $(i+1)$st through $(n/2)$th paths. These $n/2 - i$ events are identical although they are not necessarily independent. However, for the purposes of an approximation we may *assume* they are independent and compute $\mathrm{P}(D_i)$ as follows.

$$\begin{aligned}
\mathrm{P}(D_i) \;&\approx\; \prod_{j=i+1}^{n/2} \mathrm{P}(\text{paths } i \text{ and } j \text{ do not collide}) \\
&=\; (\mathrm{P}(\text{paths } i \text{ and } i+1 \text{ do not collide}))^{n/2-i} \\
&=\; (1 - \mathrm{P}(\text{paths } i \text{ and } i+1 \text{ collide}))^{n/2-i} \\
&=\; (1 - \mathrm{P}(\text{paths 1 and 2 collide}))^{n/2-i} \\
&\geq\; \left(1 - \frac{3h-1}{n-2}\right)^{n/2-i}
\end{aligned}$$

The inequality is a consequence of (19). We expect the inequality to be an approximate equality when h is small. The expected number of successful paths is then

$$\begin{aligned}
\text{Expected \# non-colliding} &= \sum_{i=1}^{n/2} \mathrm{P}(D_i) \\
&\approx \sum_{i=1}^{n/2} \left(1 - \frac{3h-1}{n-2}\right)^{n/2-i} \\
&= \frac{n-2}{3h-1}\left(1 - \left(1 - \frac{3h-1}{n-2}\right)^{n/2}\right) \qquad (20)\\
&\approx \frac{n-2}{3h-1} \qquad (21)
\end{aligned}$$

because $(1 - x/n)^{n/2} \approx e^{-x/2}$ decreases rapidly with x. This calculation, although based on an incorrect independence assumption is often useful to get an estimate of the number of non-colliding paths that we can expect to find.

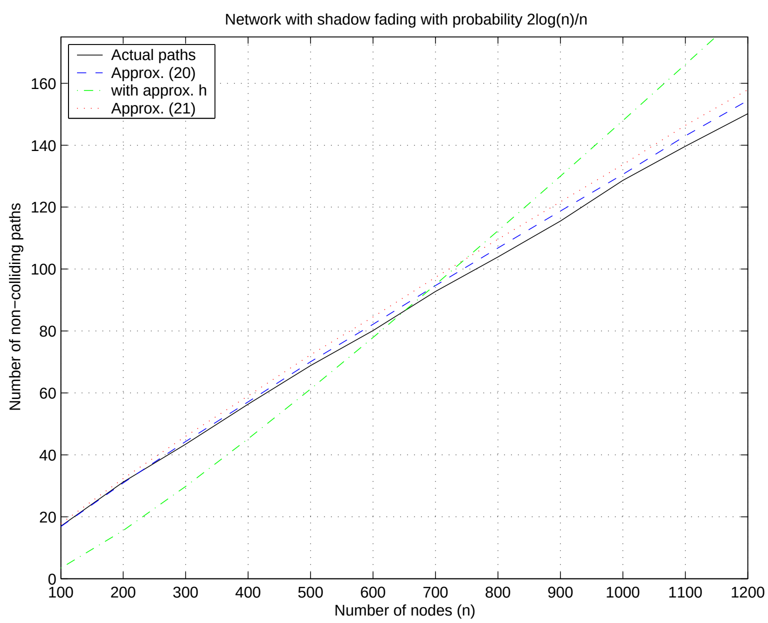

Figure 4: Number of computer-found non-colliding paths versus n for a shadow-fading model with connection probability $2(\log n)/n$ (solid curve) versus n. Also shown are the approximation (20) (dashed curve closest to solid curve) and the approximation (21) (next-closest dashed curve) using values of h obtained in the computer simulation. The dash-dotted curve is (20) computed using $h = \log(n)/\log(np)$.

We observe that in [19] vertex-disjoint paths are found successively and the nodes that are used in paths for source-destination pairs $1, \ldots, i$ are eliminated entirely from the graph before finding the path for the $(i + 1)$st pair. The paper adroitly proves that at each stage the remaining graph has edges that are "approximately" i.i.d. (from the appropriate distribution). The approximation we use above deals with the loss of the i.i.d. property by simply ignoring it. Figure 4 shows that the approximations (20) and (21) can be very accurate. The figure shows the number of computer-found non-colliding paths obtained in the shadow-fading model in Section 8.1 with link probability $p = 2(\log n)/n$. (We provide more details about this simulation in Section 10.4.) The most accurate approximation is obtained when the number of hops h in (20) and (21) is also taken from the simulation. However, we may always approximate the number of hops before the simulation as $h = \log(n)/\log(np)$. This final approximation is presented as the dash-dotted curve.

10.2 Exploiting "bad" edges to reduce interference

We use good links to establish communication paths between sources and destinations. It is similarly tempting to introduce a concept of weak or *bad* links between such paths to minimize interference. Suppose we classify as bad all links where the channel coefficients are below some threshold η_n. Those that have channel coefficients above the threshold β_n are still called good. We now wish to find non-colliding source-destination paths such that at each hop the interfering links are all bad and the communication links are all good. This is a significantly more challenging design problem than the one involving only good links for communication because every source-destination pair generates interference for the remaining $n - 2$ nodes in the network. However, we can potentially achieve higher SINRs with this approach, especially if the network has many poor connections.

10.3 Decentralized algorithm for finding non-colliding paths

Typically we would like the sources and destinations to find non-colliding paths (or the schedule) without the help of a central all-knowing being. While we expect every node to know its immediate local neighborhood, we do not expect it to know the topology of the entire network. We briefly suggest how a decentralized algorithm requiring only local information would work.

We assume that the threshold β_n is known to each node. Hence, each node knows which of its connections are good. Any decentralized algorithm can can therefore operate directly on the derived graph $G(n, p)$ of connections that exceed β_n. Decentralized shortest path algorithms in $G(n, p)$ are well known [21]. We propose a scheme in which source-destination pairs use such an algorithm to avoid collisions. Assume that the first source finds a shortest path to its destination. It then broadcasts its path to the entire network (this generally involves transmitting the information about no more than $\log n$ nodes). Then the second source finds its shortest path to its destination while avoiding the nodes that were used in the first path. Thus, the second node works on a smaller graph than the first. This process repeats until all k source-destination pairs are satisfied. While this algorithm does not require centralized knowledge, it requires some cooperation between nodes.

The analysis of this scheme involves estimating h, the lengths of paths found and k the number of source-destination pairs that can be supported. This is generally difficult since specific nodes are constantly being removed from the graph and the remaining graph no longer has the same properties as the original graph. However, the techniques used in [19] assure us that as long as the remaining graph is large enough ($n/2$ nodes remain, say) there is still enough randomness in the remaining graph to ensure that shortest paths should not exceed the diameter of the original graph.

10.4 Simulations

We revisit some of the examples analyzed in Section 8 to see how well our analytical predictions match computer-generated simulations. We begin with the shadow-fading network analyzed in Section 8.1.

Figure 5 shows the aggregate throughput and minimum SINR of a shadow-fading network as a function of the number of nodes n in a computer-generated simulation where the channel connections are chosen as in Section 8.1. The analytical results suggest that for best throughput we should choose $p = (\log n + \omega_n)/n$ for ω_n going to infinity arbitrarily slowly. We therefore choose $p = 2(\log n)/n$. The computer simulation begins by establishing a network of n connections whose channels are drawn i.i.d. according to (11). Non-colliding paths (using the method described in Section 10.1) are established and the minimum SINR obtained along the ith path, denoted $\rho_{0,i}$, is found. The quantity $\log(1 + \rho_{0,i})$ is then computed, weighted by the number of hops on path i, summed over i, and then normalized by the total number of hops contained in all paths. This gives a measure of the throughput per path, where paths that are longer (have more hops) count more heavily in the average. This throughput-per-path is then multiplied by the number of non-colliding paths and divided by the average number of hops to provide the aggregate throughput. The resulting curve appears in Figure 5 as an increasing function of n and whose y-axis is labeled on the left. The minimum SINRs obtained along the ith path $\rho_{0,i}$ are averaged over i and displayed as the decreasing curve whose y-axis is labeled on the right. As predicted in Section 8.1, the aggregate throughput grows nearly linearly. The figure shows that the average SINR per path, although decreasing with n, seems to be flattening for large n; Section 8.1 shows that the SINR should asymptotically become constant.

The following applies to all simulations described in this section: (i) Computer simulations were repeated and averaged approximately 100–200 times, depending on the size of the network and variability of the results; (ii) The nodes have unit transmit power $P = 1$ and noise variance $\sigma^2 = 0.1$. Hence, on a unit channel and in the absence of interference, the SNR is 10dB; (iii) We do not prescribe an SINR threshold. Rather, we accept any non-colliding path and use its resulting SINR in our averages. We believe this to be reasonable in practice. The threshold ρ_0 has the analytical merit that it guarantees a certain throughput; (iv) The figures often show two plots; the aggregate throughput generally given by an increasing function of n and whose scale is on the left y-axis, and the average minimum SINR generally given by a decreasing function of n and whose scale is on the right y-axis; (v) Although the analysis in the paper uses logarithms with base e, the throughputs in the figures are given in bits/channel-use.

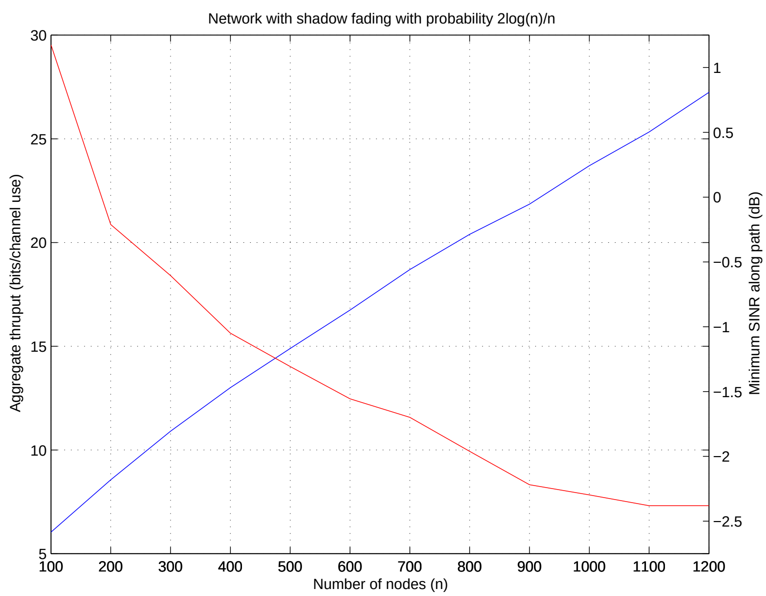

Figure 5: Aggregate throughput and minimum SINR versus number of nodes n in a shadow fading network with connection probability $p = 2(\log n)/n$. The left y-axis contains the scale for this increasing function of n. We see that the aggregate throughput increases nearly linearly. The average SINR obtained along the paths (see scale on the right y-axis) drops with n, and according to the results in Section 8.1 should asymptotically go to zero as $1/\log\log n$.

Figure 6 shows the aggregate throughput and minimum SINR of the same shadow-fading network, this time as a function of p for a fixed $n = 1000$ nodes. We see from the figure that the maximum throughput is attained when $p \approx 0.008$. Section 8.1 predicts that the maximum throughput is achieved when $p =$

$(\log n + \omega_n)/n = 0.0069 + \omega_n/n$. Ignoring the ω_n term, we see a good match between the theory and the simulation.

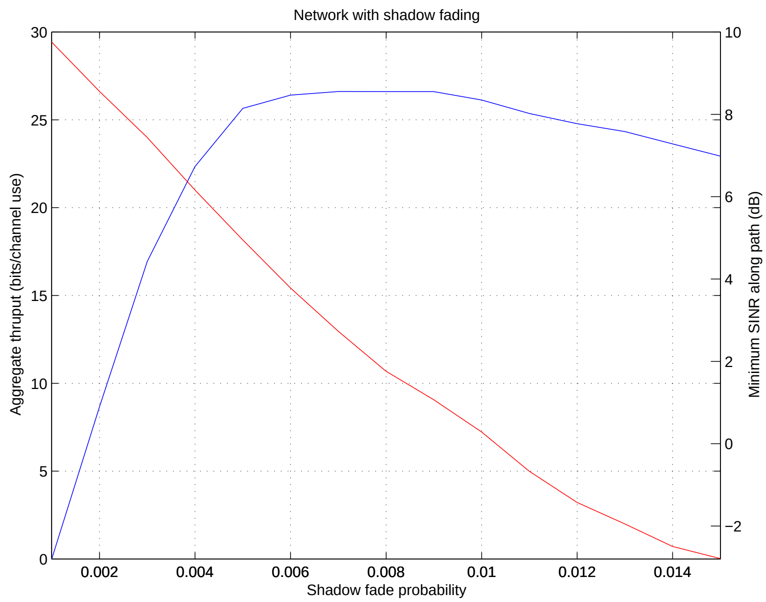

Figure 6: Aggregate throughput and minimum SINR versus connection probability p in a shadow-fading network of 1000 nodes. We see that the throughput is maximized at $p \approx 0.008$, which is not far from $(\log 1000)/1000 \approx 0.0069$, the large-$n$ maximizing p predicted in Section 8.1.

Figure 7 shows the aggregate throughput and minimum SINR of a network with exponential fading analyzed in Section 8.2 as a function of n. For large enough n the optimum threshold is $\beta = (\log n)/2$ and k should be chosen as large as possible. For purposes of illustration, we therefore choose k as large as possible, even for the relatively small values of n that we consider. (In this particular example smaller values of k can yield higher total throughput when n is small.) The result is a throughput that grows approximately logarithmically with n, as predicted theoretically. The figure also shows that choosing a β that is constant has a detrimental effect on the throughput. Similarly, choosing a β that grows faster than logarithmically would also be detrimental.

Figure 8 shows the aggregate throughput and minimum SINR of the decay-density network (as a function of n) described in Section 8.3. The parameters used in the simulation are $d = 1$, $\Delta = 1$, and $m = 3$. This is equivalent to placing nodes with unit spacing in a two-dimensional lattice and assuming a power-decay that decreases as $1/r^3$. The figure shows that the throughput grows approximately linearly, as predicted by equation (18).

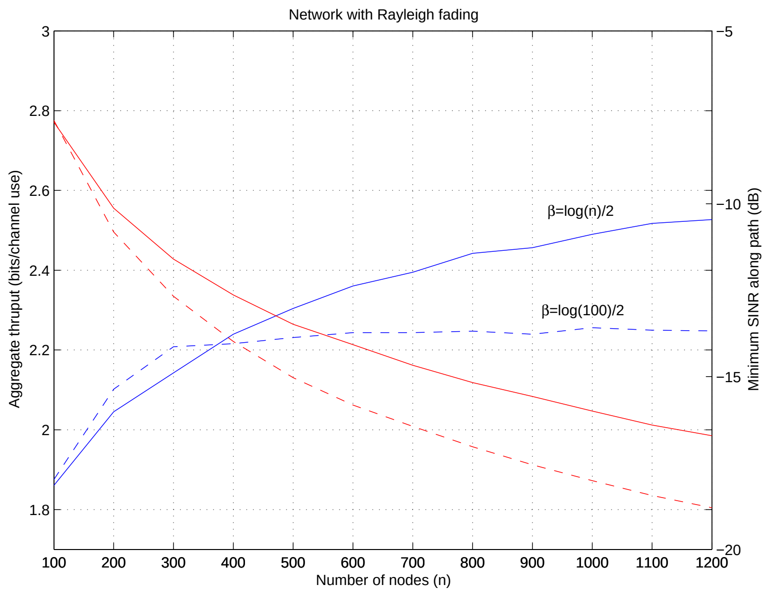

Figure 7: Aggregate throughput and minimum SINR versus number of nodes n in a network with exponential fading. We see that the throughput grows logarithmically using the optimum β computed in Section 8.2. The average SINR obtained along the paths decays approximately as $(\log n)/n$. Shown in dashed lines is the detrimental effect of choosing a constant $\beta = (\log 100)/2$.

These simulations show that Theorem 1, although designed for large n, is also accurate for finite n.

10.5 Conclusion

Our model for shared-medium wireless networks uses channels chosen according to a common distribution. We have devised a method of operating this network using relays and provided an achievable aggregate throughput as a function of the distribution. Distributions that have a certain sparsity of "good" connections seem to fare best and provide near-linear throughputs. We show that there exists an optimum amount of shadow fading that a network should have—any more or any less degrades the throughput. We hope that these results provide guidelines to the design of networks including, paradoxically, possible obstacle placement if the network is "over-connected."

We have only touched on decentralized schemes for choosing relay nodes and we have given a brief description of an upper bound on the achievable throughput. We do not generally know how sensitive our throughput results are to relaxing the i.i.d. assumption on the channel coefficients. A case where the channel coefficients are independent but have distribution that depends on distance was examined in Section 8,

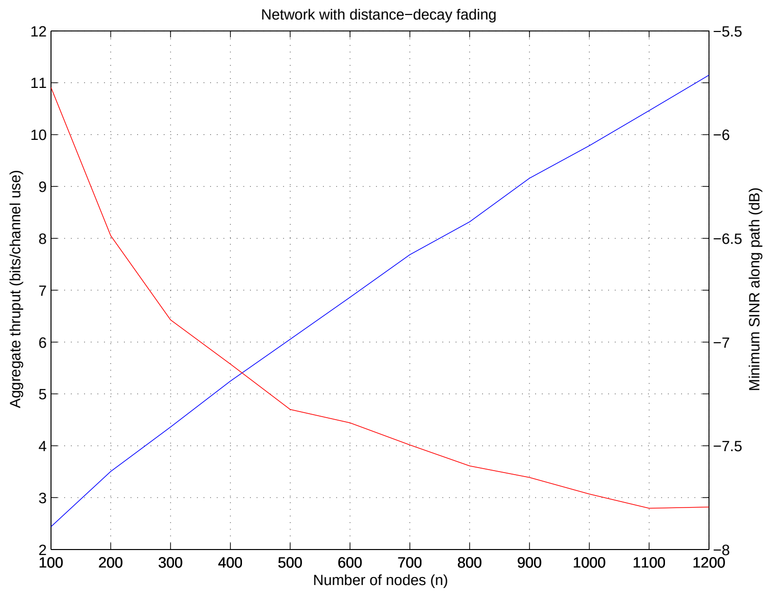

Figure 8: Aggregate throughput and minimum SINR versus number of nodes n in the decay-density network analyzed in Section 8.3. Equation (18) (for $m > 2$) predicts that the throughput should grow approximately linearly.

where we argued that at the low connection probabilities that we require the sensitivity to distance was low. It remains to be seen whether this sensitivity is low more generally.

References

[1] F. Baccelli, M. Klein, M. Lebourges and S. Zuyev, "Stochastic geometry and architecture of communication networks," *J. Telecommunication Systems*, vol. 7, pp. 209–227, 1997.

[2] B. Bollobás, *Random Graphs*, 2nd ed., Cambridge: University Press, 2001.

[3] O. Dousse and P. Thiran, "Connectivity versus capacity in dense ad hoc networks," *Proc. 23rd INFOCOM*, Hong Kong, Mar. 2004.

[4] G. J. Foschini, "Layered space-time architecture for wireless communication in a fading environment when using multi-element antennas," *Bell Labs. Tech. J.*, vol. 1, no. 2, pp. 41–59, 1996.

[5] M. Gastpar and M. Vetterli, "On the capacity of wireless networks: the relay case," *Proc. 21st INFOCOM*, New York, Jun. 2002, pp. 1577–1586.

[6] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad-hoc wireless networks," *IEEE/ACM Trans. on Networking*, vol. 10, pp. 477–486, Aug. 2002.

[7] R. Hekmat and P. van Mieghem, "Degree distribution and hopcount in wireless ad-hoc networks," *Proc. 11th IEEE Int. Conf. on Networks (ICON)*, Sydney, Australia, pp. 603–609, Sep. 2003.

[8] R. Hekmat and P. van Mieghem, "Study of connectivity in wireless ad-hoc networks with an improved radio model," *Proc. 2nd Workshop on Model. and Optim. in Mobile, Ad Hoc and Wireless Networks*, Cambridge, UK, Mar. 2004.

[9] O. Léveque and E. Telatar, "Information theoretic upper bounds on the capacity of large extended ad hoc wireless networks," *to appear in the IEEE Trans. on Info. Theory*.

[10] M. Franceschetti, O. Dousse, D. Tse and P. Thiran, "On the throughput capacity of random wireless networks," submitted to *IEEE Trans. Info. Theory*, 2004.

[11] P. Gupta and P. R. Kumar "The capacity of wireless networks," *IEEE Trans. Info. Theory*, vol. 46, pp. 388–404, Mar. 2000.

[12] P. Gupta and P. R. Kumar, "Towards an information theory of large networks: an achievable rate region," *IEEE Trans. Info. Theory*, vol. 49, pp. 1877–1894, Aug. 2003.

[13] I. E. Telatar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecom.*, vol. 10, pp. 585–595, Nov. 1999.

[14] O. Tonguz and G. Ferrari, "Connectivity and transport capacity in ad hoc wireless networks," *Proc. 32nd IEEE Comm. Theory Workshop*, Mesa, AZ, Apr. 2003.

[15] L.-L. Xie and P. R. Kumar, "A network information theory for wireless communication: Scaling laws and optimal operation," *IEEE Trans. Info. Theory*, vol. 50, pp. 748–767, May 2004.

[16] S. Toumpis and A. Goldsmith, "Capacity bounds for large wireless networks under fading and node mobility," *Proc. 41st Allerton Conf. on Comm. Cont. and Comp.*, Monticello, Oct. 2003, pp. 1369–1378.

[17] S. Weber, X. Yang, J. Andrews and G. de Veciana, "Transmission capacity of wireless ad hoc networks with outage constraints," submitted to *IEEE Trans. on Info. Theory*.

[18] F. Chung and L. Lu, "The diameter of random sparse graphs," *Advances in Appl. Math*, vol. 26, pp. 257–279, 2001.

[19] A. Z. Broder, A. M Frieze, S. Suen, E. Upfal, "An Efficient Algorithm for the Vertex-Disjoint Paths Problem in Random Graphs," *Proc 7th Symp. Discrete Algorithms*, Atlanta, 1996, pp 261– 268.

[20] T. H. Cormen, C. E Leiserson, R. L. Rivest and C. Stein *Introduction to Algorithms*, 2nd ed., The MIT Press, 2001.

[21] J. M. Abram and I. B. Rhodes, "Some shortest path algorithms with decentralized information and communication requirements," *IEEE Trans. Automatic Control*, vol. 27, pp. 570–582, 1982.

[22] M. Sharif and B. Hassibi, "On the capacity of MIMO broadcast channels with partial side information," *IEEE Trans. Info. Theory*, vol. 51, pp 506–522, Feb. 2005.