The Impact of Link Layer Model on the Capacity of a Random Ad hoc Network

Vivek P. Mhatre
Intel Research, Cambridge,
United Kingdom.
Email: mhatre@gmail.com

Catherine P. Rosenberg
Department of Electrical and Computer Engineering,
University of Waterloo, Canada.
Email: cath@ece.uwaterloo.ca

***Abstract*— The problem of determining asymptotic bounds on the capacity of a random ad hoc network is considered[1]. Previous approaches assumed a threshold-based link layer model in which a packet transmission is successful if the SINR at the receiver is greater than a fixed threshold. In reality, the mapping from SINR to packet success probability is continuous. Hence, *over each hop*, for every finite SINR, there is a non-zero probability of packet loss. With this more realistic link model, it is shown that for a broad class of routing and scheduling schemes, a *fixed fraction of hops* on each route have a *fixed non-zero packet loss probability*. In a large network, a packet travels an asymptotically large number of hops from source to destination. Consequently, it is shown that the cumulative effect of per-hop packet loss results in a per-node throughput of only $O\left(\frac{1}{n}\right)$ (instead of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$ as shown previously for the threshold-based link model).**

I. Introduction

The problem of determining the capacity of such ad-hoc networks was first systematically analyzed by Gupta and Kumar in [1]. This work has motivated several works under additional assumptions such as node mobility ([6] and references therein), network information theoretic view of network capacity [3], ultra-wideband links [8], etc.

In [1], Gupta and Kumar derive asymptotic bounds on the capacity of a random ad hoc network. In a random ad hoc network, nodes are deployed randomly and uniformly over the surface of a sphere of unit area. Each node picks a random node as its destination node, and sends packets to that node by using multi-hop communication. The authors then demonstrate that under certain simplifying link layer assumptions for successful packet reception, each node can achieve a throughput of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$ packets per second. The authors also provide a scheduling and routing strategy in order to achieve this throughput. The authors assume a link layer model in which, if the Signal to Interference and Noise Ratio (SINR) at the receiver is greater than a certain threshold β, then the packet is received successfully by the receiver with probability one.

In reality, for a given modulation and coding scheme, as long as there is some noise and interference, i.e., as long as the SINR is finite, there is always a non-zero probability of packet error, and this probability of error approaches zero as the SINR approaches infinity [2]. In other words, over *each hop*, the mapping from SINR to packet success probability is a continuous function that approaches one only asymptotically as the SINR approaches infinity. While the threshold-based packet reception model used in [1] is a reasonable choice for successful packet reception in a single hop network such as a cellular network, we argue that it needs to be refined when applied to a multi-hop network. In an adhoc network, each packet traverses *multiple hops*. The links of these hops receive interference from other ongoing transmissions which could potentially corrupt the packet transmission over the given link. When a packet is relayed over a large number of links, each of which is likely to drop the packet with a certain probability, the end-to-end throughput depends on the probability of end-to-end packet delivery. In [1], the capacity determining constraint is the number of source-destination paths that pass through a cell, i.e., the relaying burden of a cell. However we show that other factors such as interference, and the end-to-end packet error probability also have a big impact on network capacity. We show that these factors can bring down the network capacity to $O\left(\frac{1}{n}\right)$, instead of previously reported achievable bound of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$.

The paper is organized as follows. We motivate the problem, and outline our approach in Section II. We discuss some of the related work in Section III. In Section IV, we formulate and solve the capacity problem. Finally, we present our conclusion and future directions in Section V.

[1] A more detailed version of this work is currently under review in IEEE Transactions on Information Theory, [7].

II. Motivation and Approach

The arguments in this section are only meant to provide intuitive insights about the problem and our approach. We provide precisely proved results for all the arguments in Section IV.

For the routing and scheduling scheme proposed in [1] to achieve a per-node throughput of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$, it can be shown that a packet traverses $\Theta\left(\sqrt{\frac{n}{\log n}}\right)$ intermediate hops on its way from a source node to its destination node. Thus asymptotically, the number of hops that a packet has to travel from source to destination goes to infinity as n scales. The authors assume that since the SINR over each hop is at least β, every packet transmission is *always* successful. Let us call this the ideal link model.

Consider another link model in which every link in the network is reliable with a probability $p < 1$, i.e., each packet transmitted on the link is received successfully at the receiver with a probability of p. Let us call this the probabilistic lossy link model. For simplicity, assume that there are no retransmissions (this aspect is considered later in subsection IV-B). Since there are $\Theta\left(\sqrt{\frac{n}{\log n}}\right)$ hops from source to destination, the probability that the packet reaches its destination scales as $p^{\sqrt{\frac{n}{\log n}}}$. This quantity is $O\left(\sqrt{\frac{\log n}{n}}\right)$, since p^m is $O\left(\frac{1}{m}\right)$ as m tends to infinity. Hence, under the probabilistic lossy link model, the achievable throughput is $O\left(\frac{1}{\sqrt{n \log n}} \times \sqrt{\frac{\log n}{n}}\right)$. This results in an end-to-end throughput of $O\left(\frac{1}{n}\right)$ for the probabilistic lossy link model.

A realistic link model lies somewhere between the ideal link model and the probabilistic lossy link model. In a realistic link model, the probability of link reliability, p, is a continuous function of the SINR of the link. The SINR in turn depends on factors such as transmitter-receiver separation, power of interference from simultaneous transmissions, and noise power. These factors are different for different links along a path, and hence packets are lost with different probabilities along different links depending on the SINR of the link. We show that even under such a link layer model, due to packet losses along the intermediate links, the end to end throughput scales as $O\left(\frac{1}{n}\right)$ for a broad range of routing and scheduling schemes.

III. Related Work

Throughout this paper, we refer to the work of Gupta and Kumar on the capacity of random ad hoc networks [1]. In this work, the authors assume a simplified link layer model in which each packet reception is successful if the receiver has an SINR of at least β. The authors assume that each packet is decoded at every hop along the path from source to destination. No co-operative communication strategy is used, and interference signal from other simultaneous transmissions is treated just like noise. For this communication model, the authors propose a routing and scheduling strategy, and show that a per-node throughput of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$ can be achieved.

In [3], the authors discuss the limitations of the work in [1], by taking a network information theoretic approach. The authors discuss how several co-operative strategies such as interference cancellation, network coding, etc. could be used to improve the throughput. However under decode-and-forward communication model, the throughput of a random ad hoc networks is shown to be $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$. In [1], the authors consider another model of ad hoc networks called arbitrary networks, and show that a per-node throughput of $\Theta\left(\frac{1}{\sqrt{n}}\right)$ is achievable. In [4], the authors propose using a backbone-based relaying scheme that uses power control and achieves the $\Theta\left(\frac{1}{\sqrt{n}}\right)$ throughput bound even for random networks. The latter approach has also been used in [5].

However, just as in [1], all the above mentioned works assume that over each link a certain non-zero rate can be achieved. They do not take into account the fact that in reality, such a rate is achieved *with a probability of bit error arbitrarily close (but not equal) to zero*. Once the coding and modulation scheme is fixed, the function corresponding to the probability of bit error is also fixed.

IV. Problem Formulation and Solution

We use the same terminology for $\Theta(.), O(.), \Omega(.)$, etc. as used by Gupta and Kumar in [1]. Consider a sphere of unit area, say S^2, over which n nodes are deployed randomly and uniformly. Each node picks a random node which is its destination node, and sends packets to this destination node. The authors in [1] use the following criterion for successful packet reception. If the SINR at a receiver is greater than a certain fixed threshold β, then the packet is successfully decoded by the receiver. In other words, a transmission from node i to node j is successful if:

$$\frac{\frac{P}{|X_i - X_j|^\alpha}}{N + \sum\limits_{k \in T,\ k \neq i} \frac{P}{|X_k - X_j|^\alpha}} \geq \beta, \tag{1}$$

where T is the set of all the nodes that are transmitting simultaneously with node i, and $\alpha > 2$ is the propagation loss exponent.

In our analysis, to begin with, we make the following assumptions:

1) The region S^2 is partitioned into a Voronoi tessellation such that each cell contains a circle of radius ρ_n, and each cell is contained inside a circle of radius $2\rho_n$. We assume that ρ_n is at least $\Omega\left(\sqrt{\frac{\log n}{n}}\right)$ to ensure network connectivity with high probability [1]. To ensure that the cell sizes shrink as n scales, $\rho_n \to 0$ as n approaches infinity. Each node transmits at a fixed power level P.
2) For a given modulation and coding scheme, the probability of a successful packet reception is a continuous increasing function of SINR that approaches one as the SINR approaches infinity. Hence, for every finite SINR value, we assume that there is a non-zero probability of packet loss. Even as n scales, the modulation and coding scheme is fixed. We also assume that over each hop, if a packet is not received successfully by the receiver, the transmitter does not retransmit the packet, and the packet is lost. We comment on how retransmission strategy can be taken into account in our analysis in subsection IV-B.
3) A scheduling algorithm that guarantees each cell a transmission opportunity at least once every K slots, where K is a finite number that is independent of n. In other words, the length of the schedule is bounded even as n scales. In [1], it was shown that such a scheduling strategy *exists*, and it guarantees an SINR of at least β for all the scheduled transmissions.
4) A routing scheme in which packets are routed along straight line paths between source-destination pairs, i.e., every cell that intersects the straight line joining a

source-destination pair, relays the packets of that pair (the routing policy used in [1]).

The key assumption that distinguishes our work from the previous approaches is assumption (2). The remaining three assumptions are also used by Gupta and Kumar in [1]. We later on relax some of the above assumptions.

A. Per-node Capacity under Assumptions (1) to (4)

Consider Fig. 1. Let L_i be the line segment along the surface of S^2 that connects the ith source-destination pair (henceforth referred to as the ith connection). We also use L_i to denote the length of the line segment joining the ith source-destination pair. As per assumption (4), the packets of the ith connection are relayed hop-by-hop by every cell which intersects line L_i. Over each hop, any node in the relaying cell may forward the packet. The scheduling algorithm, and the uniform cell sizes ensure that communication between any two nodes in the neighboring cells is possible by guaranteeing that the SINR at the receiving node is greater than or equal to β (see [1] for more details). We require the following three lemmas to obtain the capacity of a random ad hoc network under assumptions (1) to (4).

Lemma 1 shows that the number of hops from source to destination in bounded asymptotically.

Lemma 1: For the routing scheme in assumption (4), the number of hops H_i for connection i is $\Theta\left(\frac{L_i}{\rho_n}\right)$. More precisely,

$$\frac{1}{8}\frac{L_i}{\rho_n} \leq H_i \leq \frac{16}{\pi}\frac{L_i}{\rho_n}.$$

Proof: Since, each Voronoi cell is contained in a circle of radius $2\rho_n$, the maximum distance between a point in a given cell, and a point in its neighboring cell is $8\rho_n$ (see Fig. 1). Thus, over each hop, the maximum distance that a packet can cover is upper bounded by $8\rho_n$. Hence the lower bound. For the upper bound, if we look at a strip that is $4\rho_n$ wide on either sides of L_i (see Fig. 1), we observe that if a cell is used as a relay cell, then it must lie *entirely* within this strip. Using this observation, and the fact that the size of each cell is lower bounded, we can upper bound the number of hops (see [7] for details). ∎

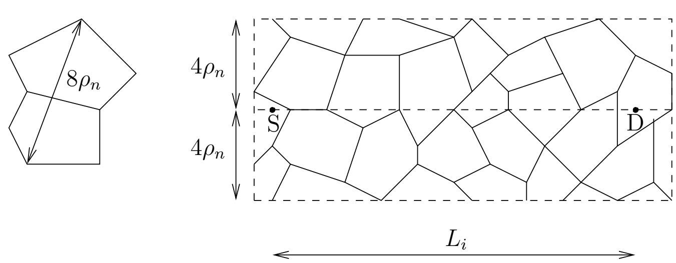

Fig. 1. Bounding the number of hops H_i for connection i.

The following lemma shows that not all hops of a connection are arbitrarily short. Or equivalently, at least a fixed fraction of hops of each connection are longer than a certain distance. This in term implies that the power received by the receivers of these hops is upper bounded.

Lemma 2: Fix t such that $0 < t < 1$. For connection i, out of H_i total hops, let h_i hops be such that each of these hops covers a distance of less than $t\rho_n$. Then

$$H_i - h_i \geq \frac{L_i}{\rho_n}\left(\frac{1-\frac{16t}{\pi}}{8-t}\right)$$

Thus, for the above $H_i - h_i$ hops, the signal received at the receiver is at the most $P(t\rho_n)^{-\alpha}$, where P is the transmit power common to all the nodes, and α is the propagation loss exponent ($\alpha > 2$).

Proof: Since h_i hops each cover less than $t\rho_n$ distance, the leftover distance, which is at least $L_i - h_i t\rho_n$ has to be covered by the rest of the $H_i - h_i$ hops. Each of these $H_i - h_i$ hops can cover a distance of at the most $8\rho_n$. This gives the desired lower bound on the number of hops that are longer than $t\rho_n$. Clearly, for these hops, the received power is upper bounded by $P(t\rho_n)^{-\alpha}$. The details of the proof have been omitted due to space constraints. See [7] for details. ∎

Having shown that the received power is upper bounded, we now show that the interference is lower bounded for these hops. These two claims together imply that the SINR of these hops is upper bounded, and hence there is a fixed non-zero packet loss over these links.

In [1], it was shown that there exists a scheduling policy that ensures an SINR of at least β at the receiver of every scheduled transmission. This scheduling policy corresponds to a graph coloring problem, and it was shown in [1] that the maximum number of colors required to color all the cells is upper bounded by $1 + c_1$, where c_1 is a fixed constant that is independent of n (see Lemma 4.4 in [1]). Using this scheduling scheme, each cell gets a transmission opportunity at least once every $1 + c_1$ time slots. Thus, with respect to assumption (3), $K = 1 + c_1$.

The objective of the following lemma is to show that except for a small fraction of hops, all the remaining hops of a connection receive a certain minimum amount of interference from other ongoing transmissions.

Lemma 3: Fix $M > 9$. Let N_i be the number of hops of connection i such that there is no simultaneous interfering transmission within a circle of radius $(M + 8)\rho_n$ around the receivers of those hops. To avoid tedious boundary conditions, let us not count the hops containing the source and the destination nodes in N_i. Then,

$$N_i \leq \frac{L_i}{\rho_n}\left(\frac{2(1+c_1)}{M}\right),$$

where c_1 is the constant from Lemma 4.4 in [1].

Proof: The proof exploits the fact that several co-channel transmit-receive pairs are within $M\rho_n$ distance of each other for a fixed M. Since the proof involves a lot of technical details, due to space constraints, we direct the reader to [7] for the proof. ∎

Let A_i be the set of hops of connection i for which the received signal is at the most $P(t\rho_n)^{-\alpha}$. Then, given $\epsilon_1 > 0$, we can find $t > 0$ small enough so that using Lemma 2,

$$|A_i| \geq \frac{L_i}{\rho_n}\left(\frac{1}{8} - \epsilon_1\right), \tag{2}$$

where $|A_i|$ is the number of elements in set A_i. Let B_i be the set of hops of connection i for which there is no simultaneous transmissions within a distance of $(M+8)\rho_n$ of the receiver. Using Lemma 3, $|B_i| = N_i$, and given $\epsilon_2 > 0$, we can choose M large enough so that

$$|B_i| \leq \frac{L_i}{\rho_n}\epsilon_2 \tag{3}$$

Thus using (2) and (3),

$$\begin{aligned}|A_i \cap {B_i}^c| &\geq |A_i| - |B_i| \\ &\geq \frac{L_i}{\rho_n}\left(\frac{1}{8} - \epsilon_1\right) - \frac{L_i}{\rho_n}\epsilon_2\end{aligned}$$

If we pick $\epsilon_1 = \epsilon_2 = 1/32$, and choose $t = t_0$ and $M = M_0$ corresponding to this choice of ϵ_1, ϵ_2, then

$$|A_i \cap {B_i}^c| \geq \frac{L_i}{16\rho_n} \tag{4}$$

Note that $A_i \cap {B_i}^c$ is the set of hops over which the received signal is no more than $P(t_0\rho_n)^{-\alpha}$, *and* there is at least one simultaneous transmitter within a distance of $(M_0+8)\rho_n$ of the receiver. This in turn means that for these hops, the SINR is upper bounded by

$$SINR \leq \left(\frac{M_0+8}{t_0}\right)^{\alpha} = \beta_0 \tag{5}$$

Thus we have proved the following Proposition.

Proposition 1: There exist fixed constants t_0 and M_0, that do not depend on n, such that for at least $L_i/16\rho_n$ hops of connection i, the SINR is less than a fixed constant β_0 given by (5). As per assumption (2), since the SINR is upper bounded by a fixed constant β_0, the probability of successful packet reception is also upper bounded by a fixed constant $\phi(\beta_0) < 1$, where $\phi(.)$ is the mapping between the SINR and the packet success probability, and this mapping is fixed for the choice of modulation and coding.

■

Let us assume that time is slotted such that each slot is long enough to transmit a single packet of fixed length. Let $\lambda_n(i)$ be the rate in packets/slot at which source node i injects packets in the network. Although the source node injects packets at a rate of $\lambda_n(i)$ packets per slot, not all the packets make it to the destination node. The actual end-to-end throughput of connection i, denoted by $\Lambda_n(i)$ is given by,

$$\begin{aligned}\Lambda_n(i) &= \lambda_n(i)\,\text{Prob}\,\{\text{packet is received successfully over} \\ &\qquad \text{all the hops of connection } i\} \\ &= \lambda_n(i)\,\Pi_{j=1}^{H_i}\text{Prob}\,\{\text{packet is received successfully} \\ &\qquad \text{over the } j\text{th hop of connection } i\},\end{aligned}$$

where we assume that the interference signal, and noise observed by a packet at each of the hops are independent. We know from Proposition 1 that, among the H_i hops of connection i, at least $L_i/16\rho_n$ hops have a probability of packet success of no more than $\phi(\beta_0) < 1$. Thus,

$$\Lambda_n(i) \leq \lambda_n(i)\ \{\phi(\beta_0)\}^{\frac{L_i}{16\rho_n}} \tag{6}$$

Note that L_i are i.i.d. random variables. Hence if we remove the conditioning on L_i by taking expectation with respect to L_i, the end-to-end throughput Λ_n is,

$$\begin{aligned}\Lambda_n &= \mathbf{E}_L[\Lambda_n(i)] \\ &\leq \lambda_n(i)\,\mathbf{E}_L\left\{\left(\{\phi(\beta_0)\}^{\frac{1}{16\rho_n}}\right)^{L_i}\right\} \\ &= \lambda_n(i)\,\mathbf{E}_L\left[\delta^{L_i}\right]\end{aligned} \tag{7}$$

where we have substituted $\delta = \{\phi(\beta_0)\}^{\frac{1}{16\rho_n}}$. Note that in determining the average end-to-end throughput Λ_n, we take expectations at two levels; once to take into account the randomness due to the possibility of packet error on each link, and once to take into account the randomness due to the locations of the source and destination nodes. Also note that $0 < \delta < 1$. Since L_i is a line connecting two points picked at random on the surface of S^2, we can show that (see [7] for details).

$$\mathbf{E}_L\left[\delta^{L_i}\right] = \frac{2\pi\left(1+\delta^{\frac{\sqrt{\pi}}{2}}\right)}{4\pi + (\log\delta)^2} < \frac{4\pi}{(\log\delta)^2} \tag{8}$$

Using (7) and (8), and substituting $\delta = \{\phi(\beta_0)\}^{\frac{1}{16\rho_n}}$,

$$\Lambda_n < \lambda_n(i)\,\frac{1024\pi{\rho_n}^2}{(\log\phi(\beta_0))^2} \tag{9}$$

So far, i.e., in proving Lemmas 1, 2 and 3, Proposition 1, and (9) we have not made any assumptions about the exact form of ρ_n. We have just assumed (1) to (4). Let us now study the per-node throughput bound for for the choice of ρ_n in [1], i.e., when ρ_n is chosen to be the radius of a disk of area $\frac{100\log n}{n}$. In [1], it was shown through Lemma 4.8, that for this choice of ρ_n each Voronoi cell contains at least one node with high probability, and the network is connected. More precisely, there exists a sequence $\delta_n \to 0$, such that

$$\begin{aligned}&\text{Prob}\,\{\text{Number of nodes in cell } V \geq 50\log n,\ \text{for every} \\ &\qquad \text{cell } V \text{ in the tessellation}\} > 1 - \delta_n\end{aligned}$$

For the above choice of ρ_n, we have the following lemma.

Lemma 4: Assume that ρ_n is chosen as the radius of a circle of area $\frac{100\log n}{n}$. If $\lambda_n(i)$ is the rate in packets/slot at which *every* source node injects packets in the network, then with high probability,

$$\lambda_n(i) \leq \frac{1}{50\log n}$$

Proof: The scheduling algorithm proposed in [1] guarantees each cell a time slot at least once every $1 + c_1$ slots. However, since each cell contains at least $50\log n$ nodes with high probability, even if each node were to transmit only its own packets, and not relay packets, it would still get no more than one transmission opportunity every $50 \log n$ slots. Hence, the rate at which a node injects its own packets into the network can never be more than $1/50 \log n$. ∎

Using the above Lemma to bound $\lambda_n(i)$ in (9) we get

$$\Lambda_n < \frac{1}{50 \log n} \frac{1024\pi {\rho_n}^2}{(\log \phi(\beta_0))^2} \tag{10}$$

Also, since ρ_n is the radius of a circle of area $\frac{100 \log n}{n}$ on S^2,

$$\frac{\pi {\rho_n}^2}{2} < \frac{100 \log n}{n} \quad \text{(see [7] for details.)}$$

Substituting the above in (10),

$$\Lambda_n < \frac{2^{12}}{(\log \phi(\beta_0))^2} \frac{1}{n} \tag{11}$$

Thus we have proved the following proposition which is the main result of this paper.

Proposition 2: Under assumptions (1) to (4), and with ρ_n chosen to be the radius of a circle of area $\frac{100 \log n}{n}$ on S^2 (as in [1]), the per-node throughput that can be achieved, Λ_n, is $O(\frac{1}{n})$ instead of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$. ∎

The upper bound of $O\left(\frac{1}{n}\right)$ on per-node throughput can also be extended to any other choice of ρ_n that satisfies assumption (1). These results can also be generalized to the case of arbitrary routing, i.e., if assumption (3) in Section IV is relaxed. These general cases are dealt with in [7]. Thus, under a realistic link layer model, the per-node throughput of a random ad hoc network is $O\left(\frac{1}{n}\right)$ instead of $\Theta\left(\frac{1}{\sqrt{n \log n}}\right)$.

B. A Remark about Retransmission Strategy

In our analysis so far, we have assumed that over each hop, a single attempt is made to transmit the packet. This is because the original Gupta-Kumar work in [1] did not take into account retransmissions at the link layer. If a packet transmission fails, then the node does not make a retransmission attempt. In reality, a multi-hop link layer employs a hop-by-hop acknowledgment and retransmission strategy in which the transmitter retransmits the packet if its previous transmission fails. If up to R retransmissions are allowed over each hop, and if p is the probability of packet success over each transmission, then the probability that the packet will be transmitted successfully within R attempts is simply given by $1 - (1-p)^R$. This quantity is strictly less than one for finite R. Thus, even in this case, there is an upper bound on the packet success probability which is strictly less than one. Consequently, all the arguments in this paper still go through with slight modifications. We can now think of replacing p as the packet success probability by $1-(1-p)^R$, which is still a constant independent of n, and strictly less than one. As a result, the per-node throughput still remains $O\left(\frac{1}{n}\right)$.

The impact of a retransmission scheme that scales with n, on the network capacity is an interesting problem that will be addressed in an extension of this work. This problem is not easy to analyze, since unlike the work in [1], we cannot merely count the number of simultaneous packet transmissions in the network to determine capacity. We must also track packet losses from source to destination.

V. Conclusion and Extensions

The key observation in this paper is that for a broad range of routing and scheduling schemes (including the one proposed in [1]), the number of hops of each connection scales to infinity with n. In a realistic link layer model, for a given modulation and coding scheme, the probability of packet loss over any given link is non-zero for finite SINR values. With the above link layer model, we show that for a broad class of routing and scheduling policies we cannot achieve a per-node throughput of more than $O\left(\frac{1}{n}\right)$ due to the cumulative probability of packet loss over all the hops of the connection.

Our results show that packet loss due to interference, and large number of hops are the key parameters that may determine the network throughput. These parameters *do not* require any additional assumptions on the communication paradigm such as any-to-any communication (as in [1]), many-to-one communication (as in sensor networks), etc. Thus our results can be generalized to more generic network architectures such as [4], [5], and this will be addressed in an extension of this work.

Acknowledgments

This work was supported in part by the Indiana Twenty First Century Fund through the Indiana Center for Wireless Communication and Networking, and by the National Science Foundation Grant No. 0087266. We would like to thank Prof. Ravi Mazumdar (University of Waterloo, Canada), and Prof. Ness Shroff (Purdue University, USA) for their valuable comments.

References

[1] P. Gupta and P. R. Kumar, “The capacity of wireless networks,” *IEEE Transactions on Information Theory*, vol. IT-46, no. 2, pp. 388-404, March 2000.

[2] R. G. Gallager, *Information Theory and Reliable Communication*, New York, NY, USA: Wiley, 1968.

[3] Liang-Liang Xie and P. R. Kumar, “A network information theory for wireless communication: Scaling laws and optimal operation,” *IEEE Transactions on Information Theory*, Vol. 50, No. 5, May 2004.

[4] M. Franceschetti, O. Dousse, D. Tse and P. Thiran, “On the throughput capacity of random wireless networks,” Under Revision. Available at http://fleece.ucsd.edu/∼massimo/Journal/IEEE-TIT-Capacity-Submission.pdf.

[5] O. Dousse, M. Franceschetti and P. Thiran, “Information theoretic bounds on the throughput scaling of wireless relay networks,” *IEEE INFOCOM’05*, Miami, March 2005.

[6] G. Sharma, R. Mazumdar and N. Shroff, “Delay and Capacity Tradeoffs in Mobile Ad Hoc Networks: A Global Perspective,” to appear in *Infocom 2006*, Barcelona, Spain.

[7] V. P. Mhatre and C. P. Rosenberg, “The Capacity of Random Ad hoc Networks under a Realistic Link Layer Model,” under review, *IEEE Transactions on Information Theory*, submitted in April 2005. Available at http://www.arxiv.org/pdf/cs.IT/0504099.

[8] A. Rajeswaran and R. Negi, “Capacity of power constrained ad-hoc networks,” in proc. of *Infocom 2004*, Hong Kong.