

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Optimal Scaling of Multicommodity Flows in Wireless Ad-Hoc Networks: Beyond the Gupta-Kumar Barrier

Permalink

<https://escholarship.org/uc/item/39c4616f>

Author

Garcia-Luna-Aceves, J.J.

Publication Date

2008-09-29

Peer reviewed

Optimal Scaling of Multicommodity Flows in Wireless Ad Hoc Networks: Beyond The Gupta-Kumar Barrier

Shirish Karande [§], Zheng Wang [§], Hamid Sadjadpour [§], J.J. Garcia-Luna-Aceves ^{§,*}

[§]Department of Electrical Engineering and Computer Engineering

University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA

*Palo Alto Research Center (PARC), 3333 Coyote Hill Road, Palo Alto, CA 94304, USA

Email: {karandes,wzgold,hamid,jj}@soe.ucsc.edu

Abstract

We establish a tight max-flow min-cut theorem for multi-commodity routing in random geometric graphs. We show that, as the number of nodes in the network n tends to infinity, the maximum concurrent flow (MCF) and the minimum cut-capacity scale as $\Theta(n^2 r^3(n)/k)$ for a random choice of $k \geq \Theta(n)$ source-destination pairs, where $r(n)$ is the communication range in the network. We exploit the fact, that the MCF in a random geometric graph equals the interference-free capacity of an ad-hoc network under the protocol model, to derive scaling laws for interference-constrained network capacity. We generalize all existing results reported to date by showing that the per-commodity capacity of the network scales as $\Theta(1/r(n)k)$ for the single-packet reception model suggested by Gupta and Kumar, and as $\Theta(nr(n)/k)$ for the multiple-packet reception model suggested by others. More importantly, we show that, if the nodes in the network are capable of multiple-packet transmission and reception, then it is feasible to achieve the optimal scaling of $\Theta(n^2 r^3(n)/k)$, despite the presence of interference. This result provides an improvement of $\Theta(nr^2(n))$ over the highest achieved capacity reported to date. In stark contrast to the conventional wisdom that has evolved from the Gupta-Kumar results, our results show that the capacity of ad-hoc networks can actually increase with n while the communication range tends to zero!

1 Introduction

The price, performance and form factors of sensors, processors, storage elements, and radios today are at a point that the age of viable large-scale ad-hoc networks would appear to be finally upon us. Alas, Gupta and Kumar's [8] seminal work cautions us to the contrary. They studied the capacity of wireless ad-hoc networks with n nodes when re-

ceivers are static, transmit or receive one packet at a time, and the network traffic consists of n unicast sessions. Their conclusion was a rather negative one: The capacity of ad-hoc networks does not scale with an increase in network size.

Gupta and Kumar's analysis [8] applies to the traditional view of ad hoc networking in which protocols are based on a one-to-one communication paradigm aimed at avoiding multiple access interference (MAI). However, a number of recent advances in cooperative communication and generalizations of routing are challenging the long-held view that avoiding interference is the way to maximize throughput in ad hoc networks. For example, network coding (NC) [1] generalizes routing by permitting processing of packets at intermediate nodes. In certain network configurations (refer to illustrations in [9, 19]) some nodes can utilize NC to concurrently transmit multiple packets. Many-to-one and many-to-many communication is also feasible under a variety of other cooperative techniques [3, 6, 18].

Co-operative protocols that provide performance benefits in specific network configurations need not scale well with the network size. In particular, Liu et al. [15] proved another disheartening result: NC cannot increase the throughput order of wireless ad-hoc networks for multi-pair unicast applications under half-duplex communication. However, in a recent challenge, Garcia-Luna-Aceves et al. [6] call for the realization of ad hoc networks that scale by embracing MAI through the use of multi-packet reception (MPR) at the receivers. They show that, if the nodes in the network are capable of MPR, then the order capacity of a network with n unicast sessions grows as $\Theta(r(n))$, where $r(n)$ is the communication range. This represents a gain of $\Theta(nr^2(n))$ over the throughput order of $\Theta(1/nr(n))$ reported by Gupta and Kumar.

Interestingly, the prior work on the capacity of wireless networks, which we summarize in Section 2, has focused on what is attainable with specific approaches to handle MAI.

No prior work has focused on first establishing what is the optimal capacity of a wireless network in the absence of MAI, and then determining whether that capacity is attainable when MAI is present. This is precisely the focus and overall contribution of this paper.

Section 3 presents the first contribution of this paper. We model a *random network* with n nodes, a homogeneous communication range of $r(n)$, and unicast traffic for k source-destination (S-D) pairs. In the absence of interference, such a network corresponds to a *random geometric graph* (RGG) with an edge between any two nodes separated by a distance less than $r(n)$. We define a *combinatorial interference model* based on RGGs, and use it to express all the protocol models used in the past and a model that we later use to show that the optimal capacity of a wireless network is indeed attainable. We introduce a protocol model in which nodes have the ability to decode correctly multiple packets transmitted concurrently from different nodes, and transmit concurrently multiple packets to different nodes. We refer to this as the multi-packet transmission and reception (MPTR) protocol model.

Section 4 presents the second contribution of this paper, which is the characterization of the optimal interference-free capacity of a wireless network. The task of concurrently maximizing the data-rate for k S-D pairs is an instance of the multi-commodity flow problem. Hence, the *maximum concurrent multi-commodity flow-rate* (MCF) in a RGG equals the interference free capacity (i.e., the optimal capacity) of the network. To derive upper bounds on the optimal network capacity, we use the fact that the MCF is less than the minimum capacity of a multi-commodity cut for any arbitrary graph. The max-flow min-cut theorem by Ford and Fulkerson [4] establishes that this bound is tight for a single commodity. However, in general, the min-cut does not provide a tight bound on the max-flow [14]. The bound is known to be tight only for special cases [10], and in general can exhibit a gap of $\Theta(\log n)$ [14]. We establish a tight max-flow min-cut theorem for RGGs for the first time, and show that $\Theta(n^2 r^3(n)/k)$ is a tight bound on the optimal capacity of a wireless network.

Section 5 presents our third contribution, which consists of generalizing prior results by Gupta and Kumar and by Garcia-Luna-Aceves et al., and proving that the optimal capacity of wireless networks is attainable in the presence of MAI. We utilize the max-flow min-cut theorem of Section 4 to deduce tight order bounds for the capacity of random networks under various interference models. We show that the per-commodity capacity, under the protocol model suggested in [8], exhibits a tight order bound of $\Theta(1/r(n)k)$. This result generalizes Gupta and Kumar's result to any $k \geq \Theta(n)$ S-D pairs. Similarly, we generalize Garcia-Luna-Aceves et al.'s analysis for the MPR protocol model. We show that, under the MPR model, the per-commodity ca-

capacity of the network scales as $\Theta(nr(n)/k)$, which means that it is bounded away from the optimal capacity by a factor of $\Theta(nr^2(n))$. Furthermore, the analysis in [6] implicitly assumes the existence of a tight max-flow min-cut theorem for RGGs; therefore, the results in this paper fill an important gap in their analysis.

We show that MPTR achieves the optimal capacity of $\Theta(n^2 r^3(n)/k)$. Hence, MPTR provides a gain of $\Theta(nr^2(n))$ over MPR and any previously reported feasible capacity. What is just as striking is that MPTR can achieve the dual objective of increasing capacity and decreasing the transmission range as n increases. This is in stark contrast to the commonly held view that the capacity of multihop wireless networks cannot increase as the number of nodes increases. Indeed, our results demonstrate that the capacity of ad-hoc networks can actually *increase* with n while the communication range tends to zero! Section 6 addresses the impact of our results on the design of protocols for future wireless ad hoc networks.

2 Related Work

There have been many contributions on the capacity study of wireless ad hoc networks, however due to space limitations we only mention a few of them that focus on unicasting. A number of papers have extended the results by Gupta and Kumar [8], which showed a gap between the upper and lower bounds on capacity under the physical model. Franceschetti et al. [5] closed this gap using percolation theory. Several techniques aimed at improving the capacity of wireless ad hoc networks have been analyzed. Grossglauser and Tse [7] demonstrated that a non-vanishing capacity can be attained at the price of long delivery latencies by taking advantage of long-term storage in mobile nodes. Some works have demonstrated that changing physical layer assumptions such as using multiple channels [13] or MIMO cooperation [17] can change the network capacity.

Ozgun et al. [17] proposed a hierarchical cooperation technique based on virtual MIMO to achieve linear capacity. They showed that the optimal per-session capacity of an ad-hoc network is bounded as $O(n \log n)$, and a constant per-session capacity of $\Theta(1)$ is achievable. Our work is significantly different from this work, in terms of the model and assumptions used to derive the results. Ozgun et al. consider the information-theoretic model, and assume that the network employs heterogeneous hop-sizes, at times requiring a direct communication between widely separated nodes. In contrast, our work is based on the protocol model and assumes a homogenous transmission range, which is a more realistic assumption.

Cooperation can be extended to the simultaneous transmission and reception at the various nodes in the network, which can result in significant capacity improvement [3].

As we have stated, Garcia-Luna-Aceves et al. [6] showed that using MPR at the receivers can increase the order capacity of wireless networks subject to unicast traffic.

A generalization of the max-flow min-cut theorem to multiple commodities is not feasible in arbitrary graphs. However, Leighton and Rao [14] showed that the gap between the max-flow and min-cut is at most $\Theta(\log n)$. In a recent work, Madan et. al. [16] have utilized the Leighton and Rao's work to derive bounds for the capacity of ad-hoc networks. They focus on the special case where $k = \Theta(n^2)$ and the *single packet reception* (SPR) model. Our work is inspired by the analysis of Leighton and Rao. Since we are able to deduce a tight max-flow min-cut theorem, our bounds are tighter than those reported by Madan et al. [16]. Furthermore, our work is applicable to a wider variety of protocol models and traffic patterns.

3 Preliminaries

3.1 Network Model

For a continuous region R , we use $|R|$ to denote its area. We denote the cardinality of a set S by $|S|$, and by $\|x - y\|$ the distance between nodes x and y . Whenever convenient, we utilize the indicator function $1_{\{P\}}$, which is equal to one if P is true and zero if P is false. $Pr(E)$ represents the probability of event E . We say that an event E occurs with high probability (w.h.p.) as $n \rightarrow \infty$ if $Pr(E) > (1 - (1/n))$. We employ the standard order notations O, Ω , and Θ .

We assume a random wireless network with n nodes distributed uniformly in a unit-square. In our model, as n goes to infinity, the density of the network also goes to infinity. Therefore, our analysis is applicable to dense networks. Furthermore, we assume a fixed transmission range $r(n)$ for all the nodes in the network. Thus, the network topology can be characterized using a RGG, which we denote by G_r and define next.

Definition 3.1 Random Geometric Graph G_r :

We associate a directed graph $G_r(V_r, E_r)$ with a wireless network formed by distributing n nodes uniformly in a unit square. We represent the node-set by $V = \{1, \dots, n\}$. Let the locations of these nodes be given by $\{X_1, \dots, X_n\}$, the edge-set is then $E = \{(i, j) \mid \|X_i - X_j\| \leq r(n)\}$.

While the results in this paper can be extended to undirected graphs, it is more convenient for us to use directed graphs because we use edge-coloring techniques in our work. In the case of undirected graphs, the argument should be based on vertices (receivers) rather than edges. Also note that we permit two edges for a pair of connected vertices with possibly different capacity in each direction.

We assume that the network operates using a slotted channel and, in the absence of interference, the data rate

in each time slot for every transmitter-receiver pair is a constant of value W bits/slot. Given that W does not change the order capacity, we normalize its value to 1. Hence, the interference-free capacity of each edge in G_r is equal to 1.

Gupta and Kumar [8] have proved the following criteria for the connectivity of G_r .

Lemma 3.2 For a random distribution of n nodes in a unit-square, the graph G_r is connected with high probability as $n \rightarrow \infty$, iff: $r(n) \geq r_c(n) = \Theta(\sqrt{\log n/n})$.

In a dense network, interference is the primary constraint on network capacity. Like [16], we describe the interference of a network by the following generic model.

Definition 3.3 Combinatorial Interference Model:

The interference model for the graph¹ $G(V, E)$ is determined by a function $I : E \rightarrow P(E)$, where $P(E)$ is the power set of E , i.e., the set of all possible subsets of E . For every $e \in E$, $I(e)$ represents an interference set such that, a transmission on edge e is successful if and only if (iff) there are no concurrent transmissions on any $\hat{e} \in I(e)$. An interference model can be restricted to a sub-graph $H(V_H, E_H)$ by defining a function $I_H : E_H \rightarrow P(E_H)$ such that $I_H(e) = I(e) \cap E_H$.

The various protocol models that have been proposed in the past can now be expressed as special cases on G_r .

Gupta and Kumar [8] studied a *single packet reception* (SPR) protocol model under which a transmission from node i to receiver j is successful iff $\|X_i - X_j\| \leq r(n)$ and if $\|X_j - X_k\| \geq (1 + \eta)r(n)$ for any other transmitter k . Here η is a guard-zone that is assumed to be constant for the entire network. Moreover, all the nodes operate in half-duplex mode. The following definition expresses this model in terms of the notation we have introduced.

Definition 3.4 Single-Packet Reception (SPR) Model:

Let $e = (e^+, e^-) \in E$, then the interference set for edge e is defined as

$$\begin{aligned} I_{SPR}(e) &= J(e) - e, \\ J(e) &= \{\hat{e} \in E_r \mid \|X_{\hat{e}^+} - X_{e^-}\| \leq (1 + \eta)r(n)\} \setminus \{e\} \end{aligned}$$

Garcia-Luna-Aceves et al. [6] generalized the above model to account for MPR capability at the receivers. According to their *MPR protocol model* a node i can simultaneously receive all the packets transmitted by nodes within a distance $r(n)$ iff there are no transmitters at a distance greater than $r(n)$ but less than $(1 + \eta)r(n)$. Furthermore, if a node j transmits a packet to node i , then it cannot simultaneously transmit a packet to any other node in the network.

¹Note that G_r denotes a RGG while G represents a general graph.

Definition 3.5 Multi-Packet Reception (MPR) Model:

$$\begin{aligned} I_{MPR}(e) &= I_{SPR}(e) - (A(e) - B(e)) \quad \forall e \in E \quad \text{where} \\ A(e) &= \{\hat{e} \in E_r \mid \|X_{\hat{e}^+} - X_{e^-}\| \leq r(n)\} \quad (2) \\ B(e) &= \{\hat{e} \in E_r \mid \hat{e}^+ = e^+\} \quad (3) \end{aligned}$$

We also consider the case in which nodes have MPR and MPTR capabilities, i.e., can decode multiple concurrent transmissions and can transmit concurrently multiple packets to different nodes. We capture the MPR and MPTR capabilities with a simple yet representative, *multi-packet transmission and reception* (MPTR) protocol model, which is defined below.

Definition 3.6 Multi-Packet Transmission and Reception (MPTR) Model:

$$I_{MPTR}(e) = I_{MPR}(e) - B(e) \quad \forall e \in E \quad (4)$$

It is important to highlight some of the features of the above model. The MPTR protocol model still restricts the nodes to operate in a half-duplex mode. Moreover, this model is identical to the MPR protocol model in terms of the nodes that are permitted to transmit within the vicinity of a receiver i , i.e., both models prohibit transmission from a node j such that $r(n) < \|X_i - X_j\| \leq (1 + \eta)r(n)$. Thus, the difference between the MPR and MPTR protocol models stems purely from the fact that, under the MPTR model, a node j transmitting a packet to node i can simultaneously transmit packets to other nodes in the network.

We assume that the traffic in the network is generated by unicast communication between k source-destination (S-D) pairs. We associate a rate vector $\lambda = [\lambda_1, \dots, \lambda_k]$ with these k pairs. We assume the data rate for each S-D pair to be non-zero. Hence, without loss of generality (w.l.g) the rate vector can be written as $\lambda = [fD_1, \dots, fD_k]$ where $f \in \mathbb{R}_+$ and $D_i \in [1/2, 1]$ for $1 \leq i \leq k$. We refer to the parameter f as the *concurrent flow rate* and to $D = [D_1, \dots, D_k]$ as the *demand vector*.

Definition 3.7 Feasible Flow Rate:

Given k S-D pairs $\{(s(1), d(1)), \dots, (s(k), d(k))\}$, a rate vector $\lambda = [fD_1, \dots, fD_k]$ is feasible if there exists a spatial and temporal scheme for scheduling transmissions such that by operating the network in a multi-hop fashion, and buffering at intermediate nodes when awaiting transmission, every source $s(i)$ can send λ_i bits/sec on average to the chosen destination $d(i)$. A flow rate f is feasible for a demand vector $D = [D_1, \dots, D_k]$ iff $\lambda = [fD_1, \dots, fD_k]$ is a feasible rate vector.

Definition 3.8 Capacity of Random Networks:

The capacity per commodity of a network is $\Theta(f(n))$ if under a random placement of n nodes, a random choice of k

S-D pairs and for an arbitrary demand vector we have:

$$\lim_{n \rightarrow \infty} \Pr(cf(n) \text{ is feasible flow rate}) = 1 \quad (5)$$

$$\liminf_{n \rightarrow \infty} \Pr(c'f(n) \text{ is infeasible flow rate}) < 1 \quad (6)$$

for some $c > 0$ and $c < c' < +\infty$.

In the following sections we repeatedly utilize the well known Chernoff bounds:

Lemma 3.9 Chernoff Bounds: Consider N i.i.d random variables $Y_i \in \{0, 1\}$ with $p = \Pr(Y_i = 1)$. Let $Y = \sum_{i=1}^N Y_i$. Then for every $c > 0$ there exist $0 < \delta_1 < 1$ and $\delta_2 > 0$ such that

$$\Pr(Y \leq (1 - \delta_1)Np) < e^{-cNp} \quad (7)$$

$$\Pr(Y \geq (1 + \delta_2)np) < e^{-cNp} \quad (8)$$

3.2 Graph Theory Results

We review some definitions from graph theory. In particular, note that the task of identifying a feasible flow rate can be posed as a multi-commodity flow problem, specifically the k -commodity flow problem.

Definition 3.10 k -Commodity Flow Problem:

Consider a directed graph $G(V, E)$ with a capacity function $c : E \rightarrow [0, 1]$. Let $\{(s(1), d(1)), \dots, (s(k), d(k))\}$ be k S-D pairs, with a demand vector $D \in [1/2, 1]^k$. Let $f \in \mathbb{R}_+$ be a concurrent flow rate. Find flow functions $f_i : E \rightarrow \mathbb{R}_+$ for $1 \leq i \leq k$, which satisfy the following flow constraints: Capacity Constraint : $\forall e \in E$

$$\sum_{1 \leq i \leq k} f_i(e) \leq c(e) \quad (9)$$

Flow Conservation : $\forall v \neq s(i), d(i)$

$$\sum_{e:e^+=v} f_i(e) = \sum_{e:e^-=v} f_i(e) \quad (10)$$

Demand Satisfaction : $1 \leq i \leq k$

$$\begin{aligned} fD_i &= \sum_{e:e^+=s(i)} f_i(e) - \sum_{e:e^-=s(i)} f_i(e) \\ &= \sum_{e:e^-=d(i)} f_i(e) - \sum_{e:e^+=d(i)} f_i(e) \quad (11) \end{aligned}$$

Flow functions that satisfy the above constraints are called feasible. Other inputs to the problem being fixed, a flow rate f is said to be feasible iff the above problem has a solution. Furthermore, let f^* be the MCF such that the above problem has a feasible solution. A wireless network can be represented by an equivalent graph with capacity functions determined by the interference. Thus, the MCF in

an equivalent graph can be perceived as the maximum flow that can be routed in a network. Additionally, if w.h.p. f^* is the MCF for any graph formed by a random distribution of nodes, sources and destinations, then the capacity of the wireless network is also f^* .

Consider the following additional definitions.

Definition 3.11 Vertex Cut:

Given a node set V , a cut is the separation of the vertex set V into two disjoint and exhaustive sets (S, S^C) . We shall often reference a cut just by the set S .

Definition 3.12 Multi-commodity Cut Capacity:

Given a graph $G(V, E)$, a capacity function $c : E \rightarrow [0, 1]$ and a cut (S, S^C) . The multi-commodity cut capacity is defined as

$$\Upsilon_{G,S} = \frac{\sum_{e \in E} 1_{[e^+ \in S, e^- \in S^C]} c(e)}{\sum_{i: s(i) \in S, d(i) \in S^C} D_i} \quad (12)$$

Definition 3.13 Minimum Cut Capacity:

Given a graph $G(V, E)$ and capacity function $c : E \rightarrow [0, 1]$, the minimum multi-commodity cut capacity is defined as

$$\Upsilon_G = \min_{S \subset V} \Upsilon_{G,S} \quad (13)$$

It is well-known that the minimum cut-capacity provides an upper bound on the maximum flow rate.

Lemma 3.14 For any k -commodity flow $f^* \leq \Upsilon_G$

4 Optimal Capacity

We show that for RGGs, the MCF provides a tight approximation of the minimum-cut capacity. This relationship implies a tight characterization of the interference-free capacity of wireless ad-hoc networks with a homogenous transmission range. Our approach can be summarized as follows: For a particular demand vector, we provide an upper bound by showing that there exists a multi-commodity cut in G_r of order $O(g(n))$ and a lower bound by constructing a flow of order $\Omega(g(n))$ in a sub-graph $H_r \subseteq G_r$. These results along with the following Lemmas prove that the capacity of H_r and G_r has a tight bound $\Theta(g(n))$.

Lemma 4.1 A graph $G(V, E)$ and a sub-graph $H(V_H, E_H)$ satisfy the following two properties: (a) If f is a feasible flow-rate in H then f is feasible in G ; and (b) the capacity of a cut (S, S^C) in G is always greater than or equal to the capacity of the cut in H .

Proof: To prove Part (a) of the lemma, let $f_{H,i}$ for $1 \leq i \leq k$ be the flow functions associated with the feasible flow of rate f in H . Note that these flow functions satisfy




Figure 1. A bi-partitioning of the unit square

the constraints in Definition 3.10. We construct a flow in G of rate f with the following flow functions: For $1 \leq i \leq k$ let $f_{G,i} : E \rightarrow \mathbb{R}_+$ such that

$$f_{G,i}(e) = \begin{cases} f_{H,i}(e) & \text{if } e \in E_H \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

Now, if we show that the functions $f_{H,i}$ satisfy the flow constraints, then the flow rate f is feasible in G . Definition 3.10 states that $c(e) \geq 0$ for all the edges. As a result, $\forall e \in E - E_H$ the capacity constraints are satisfied trivially, given that $f_{G,i}(e) = 0$ for such edges. Furthermore, $\forall e \in E_H$ we have

$$\sum_{1 \leq i \leq k} f_{H,i}(e) = \sum_{1 \leq i \leq k} f_{G,i}(e) \leq c(e) \quad (15)$$

Therefore, $f_{G,i}$ satisfy the capacity constraint. In addition, note that the following equations hold $\forall v \in V$:

$$\begin{aligned} \sum_{e \in E: e^- = v} f_{G,i}(e) &= \sum_{e \in E_H: e^- = v} f_{G,i}(e) + \sum_{e \in E - E_H: e^- = v} f_{G,i}(e) \\ &= \sum_{e \in E_H: e^- = v} f_{H,i}(e) + 0 \end{aligned} \quad (16)$$

$$\begin{aligned} \sum_{e \in E: e^+ = v} f_{G,i}(e) &= \sum_{e \in E_H: e^+ = v} f_{G,i}(e) + \sum_{e \in E - E_H: e^+ = v} f_{G,i}(e) \\ &= \sum_{e \in E_H: e^+ = v} f_{H,i}(e) + 0 \end{aligned} \quad (17)$$

Eqs. (16) and (17) imply that the net in-flow and the net out-flow, under $f_{G,i}$ and $f_{H,i}$, is identical for all nodes and commodities. Therefore, $f_{G,i}$ satisfies the flow conservation and demand constraints.

To show Part (b), observe that

$$\Upsilon_{G,S} = \Upsilon_{H,S} + \frac{\sum_{e \in E - E_H} 1_{[e^+ \in S, e^- \in S^C]} c(e)}{\sum_{i: s(i) \in S, d(i) \in S^C} D_i} \quad (18)$$

Since $c(e) \geq 0$ for all edges, we have $\Upsilon_{G,S} \geq \Upsilon_{H,S}$. ■

4.1 Upper Bound

We utilize the following properties of G_r .

Lemma 4.2 *If $r(n) \geq r_c(n)$ then w.h.p. graph G_r is such that: (a) The minimum vertex degree $\nabla \geq \Theta(nr^2(n))$, and (b) the maximum vertex degree $\Delta \leq \Theta(nr^2(n))$.*

Proof: We first show that $\nabla \geq \Theta(nr^2(n))$. A node v in G_r is connected to all the nodes in a disk of radius $r(n)$ centered at v . The area of this disk is $\pi r^2(n)$. Given a uniformly random distribution of nodes, the probability of another node u lying within this disk is $\pi r^2(n)$. Consider a random variable $Y_{v,u} \in \{0, 1\}$ which is equal to one iff node u is connected to node v . The degree of node v can be written as $\deg(v) = \sum_{u \in V - \{v\}} Y_{v,u}$. Therefore, the Chernoff Bound (11) implies that $\forall c > 0$ there exists a $0 \leq \delta \leq 1$ such that

$$\Pr(\deg(v) \leq (1 - \delta)n\pi r^2(n)) < e^{-cn\pi r^2(n)} \quad (19)$$

From the union bound we obtain

$$\begin{aligned} & \Pr(\nabla \leq (1 - \delta)n\pi r^2(n)) \\ & < n\Pr(\deg(v) \leq (1 - \delta)n\pi r^2(n)) \end{aligned} \quad (20)$$

Given that $r(n) \geq r_c(n)$ we have $r(n) \geq c_1 \sqrt{\log n/n}$ for some $c_1 > 0$. Therefore, Eqs. (19) and (20) imply that

$$\Pr(\nabla \leq (1 - \delta)n\pi r^2(n)) < ne^{-c\pi c_1 \log n} = 1/n^{c\pi c_1 - 1} \quad (21)$$

Now, Lemma 3.9 tell us that we can choose $c \geq 2/(\pi c_1)$ and corresponding $0 < \delta_1 < 1$ such that

$$\Pr(\nabla \geq (1 - \delta)\pi n r^2(n)) > 1 - (1/n) \quad (22)$$

We use similar arguments to show that $\Delta \leq \Theta(nr^2(n))$ w.h.p.. This fact follows from Eq. (8), which implies that $\forall c > 0$ there exists a $0 \leq \delta \leq 1$ such that

$$\Pr(\deg(v) \geq (1 + \delta_1)n\pi r^2(n)) < e^{-cn\pi r^2(n)} \quad (23)$$

The union bound and the fact that $r(n) \geq c_2 \sqrt{\log n/n}$ implies that

$$\Pr(\Delta \geq (1 + \delta)\pi n r^2(n)) < ne^{-c\pi c_2 \log n} \quad (24)$$

Therefore, there exists a $\delta_2 > 0$ s.t

$$\Pr(\Delta \leq (1 + \delta_2)\pi n r^2(n)) > 1 - (1/n) \quad (25)$$

Consider the cut S described by Figure 1. The cut consists of all the nodes in the rectangular region of a constant area. ■

Lemma 4.3 *If the network consists of $k \geq \Theta(\log n)$ S-D pairs then w.h.p. a region R of constant area $|R|$ contains $\Theta(k)$ sources with destinations outside region R .*

Proof: Let $Y_i \in 0, 1$ be a random variable that is equal to one iff the i th S-D pair is such that $s(i)$ belongs to region R and $d(i)$ does not. Under a uniformly random placement of nodes

$$\Pr(Y_i = 1) = |R|(1 - |R|) \quad (26)$$

The total number of S-D pairs satisfying the required condition can be represented by $Y = \sum_{i=1}^k Y_i$. If $k \geq c \log n$ the Chernoff Bounds imply the existence of constants $c_1 = 1/(c|R|(1 - |R|))$, $0 \leq \delta_1 \leq 1$ and $\delta_2 > 0$ s.t.

$$\begin{aligned} \Pr(Y \geq (1 - \delta_1)|R|(1 - |R|)k) & > 1 - e^{-c_1 k |R|(1 - |R|)} \\ & \geq 1 - e^{-c_1 c |R|(1 - |R|) \log n} = 1 - (1/n) \end{aligned} \quad (27)$$

$$\begin{aligned} \Pr(Y \leq (1 + \delta_2)|R|(1 - |R|)k) & > 1 - e^{-c_1 k |R|(1 - |R|)} \\ & = 1 - e^{-c_1 c |R|(1 - |R|) \log n} = 1 - (1/n) \end{aligned} \quad (28)$$

Furthermore, consider the subset A in S defined by a strip of dimension $1 \times r(n)$. The total number of vertices in A is $\Theta(nr(n))$ because of the uniform distribution of nodes in the network. ■

Theorem 4.4 *If $r(n) \geq r_c(n)$ and $k \geq \Theta(\log n)$, then the capacity of the cut S is $\Upsilon_{G_r, S} = O(n^2 r^3(n)/k)$ w.h.p.*

Proof: According to the definition of G_r two nodes are connected iff they are separated by a distance less than $r(n)$. Consequently, if an edge cuts across S then it has to be incident upon a node at a distance less than $r(n)$ from the boundary separating S and S^C , i.e. the head of the edge should lie in the subset A of dimension $r(n) \times 1$. Furthermore, for each node in A the maximum number of edges cutting across the cut is bounded by Δ

$$\sum_{e \in E} 1_{[e^+ \in S, e^- \in S^C]} \leq |A|\Delta \quad (29)$$

In the absence of interference $c(e) = 1$ for all the edges. Hence,

$$\begin{aligned} \Upsilon_{G_r, S} & = \frac{\sum_{e \in E_r} 1_{[e^+ \in S, e^- \in S^C]}}{\sum_{i: s(i) \in S, d(i) \in S^C} D_i} \\ & \leq \frac{|A|\Delta}{\sum_{i: s(i) \in S, d(i) \in S^C} D_i} \end{aligned} \quad (30)$$

According to Lemma 4.3, there exists $c_1 > 0$ s.t. the total number S-D pairs across the cut is $c_1 k$. Furthermore, by Definition 3.10 the demand for each pair is at least $1/2$. Hence,

$$\Upsilon_{G_r, S} \leq \frac{2|A|\Delta}{c_1 k} \quad (31)$$

Lemma 4.2 implies that there exists a $c_2 > 0$ s.t. $\Delta < c_2 n r^2(n)$, while uniform distribution of nodes in the network implies that there exists a $c_3 > 0$ s.t. $|A| \leq c_3 n r(n)$. Hence,

$$\Upsilon_{G_r, S} \leq 2c_2 c_3 n^2 r^3(n) / (c_1 k) \quad (32)$$

Any cut in G_r has a capacity greater than the minimum cut capacity Υ_G . Consequently Theorem 4.5 implies the following Corollary.

Corollary 4.5 *If $r(n) \geq r_c(n)$ and $k \geq \Theta(\log n)$ then Υ_{G_r} is upper bounded as $O(n^2 r^3(n)/k)$.*

4.2 Lower Bound




Figure 2. Decomposition of network area into l^2 squarelets, each of side-length $r(n)/3$

To describe a capacity-achieving flow in a more generic setting, we use an important result from parallel and distributed computing. Consider a mesh of l^2 processing units with l processors in each row and column. Let each processor be a source and destination of exactly h packets. The problem of routing the hl^2 packets to their destinations is known as $h \times h$ permutation routing and can be characterized by the following result [12].

Lemma 4.6 *If in a single slot, each processor can transmit one packet each to its immediate horizontal and vertical neighbors, then an $h \times h$ permutation routing in a $l \times l$ mesh can be performed deterministically in $hl/2 + o(hl)$ steps.*

We utilize the following corollary that can be readily deduced from the above Lemma.

Corollary 4.7 *If a processor is capable of transmitting at least η packets to each of its neighbors in each slot, then an $h \times h$ permutation routing in a $l \times l$ mesh can be performed deterministically in $O(hl/\eta)$ steps.*

Now consider a sub-graph $H_r \subseteq G_r$ obtained by employing location based constraints on the edge-set. In order to describe these constraints, we first define a location dependent hash function.

Definition 4.8 Index Function ζ :

Divide the network area into l^2 squarelets [11] of side-length $a = r(n)/3$, as shown in Figure 2. Let ζ be a function that associates an index (i, j) with a squarelet in the i th column and j th row. Furthermore, the index assigned to each squarelet is associated with each vertex in the squarelet.

We obtain H_r by removing all edges, except those connecting two nodes in vertically or horizontally adjacent squarelets. We do not necessarily have to consider H_r in order to obtain a lower bound on the interference-free capacity. However, the performance bounds for H_r play an important role when we analyze interference constrained networks in Section 5.

Definition 4.9 Geographically Restricted Sub-Graph H_r : *The graph $H_r(V_r, E_{r,H})$ is a sub-graph of G_r with an identical node-set and an edge-set defined as:*

$$E_{r,H} = \{e \in E \mid \zeta(e^-) = (a, b) \Rightarrow \zeta(e^+) = (a \pm 1, b \pm 1)\} \quad (33)$$




Figure 3. Any two points in adjacent squarelets are within a distance $r(n)$ of each other. The proof follows from the fact that the chord of a circle lies within it.

Consider some of the properties of the squarelets and H_r .

Lemma 4.10 [11] *If $r(n) \geq r_c(n)$ then w.h.p. the total number of nodes in any squarelet are $\Theta(nr^2(n))$.*

Lemma 4.11 *If $r(n) \geq r_c(n)$ and $k \geq \Theta(n)$, then w.h.p. the total number of sources in any squarelet are $\Theta(kr^2(n))$ and the total destinations in any squarelet are $\Theta(kr^2(n))$.*

Proof: For $1 \leq i \leq k$ let $Y_{i,m} \in 0, 1$ be a random variable that is equal to one iff source $s(i)$ belongs to the m th squarelet. Let $Y_m = \sum_{i=1}^k Y_{i,m}$ represent the total number of sources in the squarelet. Because $r(n) \geq r_c(n)$, Eq.(11) implies

$$\Pr(Y_m \leq (1 - \delta)kr^2(n)) < e^{-(ck \log n)/n} \quad (34)$$

The total number of squarelets in a unit square is equal to

$$(3/r(n)) \times (3/r(n)) \leq c_1 n / \log n \quad (35)$$

Therefore, the union bound implies that

$$\begin{aligned} \Pr & \left(\text{min. no. of nodes in a squarelet} < \Theta(kr^2(n)) \right) \\ & \leq (\text{total no. of squarelets}) \times e^{-(ck \log n)/n} \\ & \leq (c_1 n / \log n) \times e^{-(ck \log n)/n} \\ & = c_1 / (n^{(ck/n)-1} \log n) \end{aligned} \quad (36)$$

Thus, $k \geq \Theta(n)$ guarantees the required convergence and hence each squarelet has at least $\Theta(kr^2(n))$ sources. The upper bound on the number of sources and the bounds on the number of destinations can be calculated similarly. ■

Theorem 4.12 *If $r(n) \geq r_c(n)$ and $k \geq \Theta(n)$, then w.h.p. the maximum flow rate f_H^* in H_r is at least $\Theta(n^2 r^3(n)/k)$.*

Proof: The proof follows from mapping various components of the above defined problem to the $h \times h$ permutation routing problem. Let us map each squarelet to a processor. Consequently, for the chosen size of squarelets, the network equates to a mesh of l^2 processors with $l = 3/r(n)$. Assume that each source intends to transmit D_i as the i^{th} element of the demand vector. Because $D_i \leq 1$, Lemma 4.11 implies that the total number of bits to be transmitted to and from each squarelet are at most $h \leq ckr^2(n)$. Finally, note that any two nodes in adjacent squarelets are within a distance $r(n)$. Fig. 3 provides a geometric proof for this fact; an alternative proof can be easily obtained by employing the Pythagoras theorem. In each slot, we can send one packet along each edge between two adjacent squarelets. Therefore, Lemma 4.10 implies that

$$\begin{aligned} \eta & = (\text{min. no. of edges between adjacent squarelets}) \\ & \leq (\text{min. no. of nodes per squarelet})^2 \leq c_1 n^2 r^4(n) \end{aligned}$$

Hence, the total number of slots γ required to complete the desired routing is

$$\begin{aligned} \gamma & \leq (c_2 h l / \eta) \\ & \leq c_2 \times (ckr^2(n)) \times (3/r(n)) \times (1/c_1 n^2 r^4(n)) \\ & = (3c_2 ck / c_1 n^2 r^3(n)) \end{aligned} \quad (37)$$

We can repeat the above routing periodically to guarantee a flow rate of $f = (1/\gamma) \geq \Theta(n^2 r^3(n)/k)$. By definition, the max-flow rate is greater than any other feasible flow rate, and the theorem follows. ■

Aggregating the above results we have the following conclusion.

Theorem 4.13 *If $r(n) \geq \Theta(\sqrt{\log n/n})$ and $k \geq \Theta(n)$, then the max-flow f_G^* in G_r can be approximated tightly by the min-cut capacity $\Upsilon_{G_r}^*$ in G_r . Moreover, the f_G^* and $\Upsilon_{G_r}^*$ scale as $\Theta(n^2 r^3(n)/k)$.*

Proof: Lemma 4.1 implies that $f_G^* \geq f_H^*$. Hence, the result follows from the lower bound provided by Theorem 4.12 and the upper bound provided by Corollary 4.5. ■

Scaling laws for $k = \Theta(n)$ have been given special attention in the literature. Hence, it is worth stating explicitly the above results for this scenario.

Corollary 4.14 *Consider an ad-hoc network described by a random placement of n nodes in a unit square, with $\Theta(n)$ S-D pairs and a homogenous transmission range of $r(n) \geq \Theta(\sqrt{\log n/n})$. The interference-free capacity of the network scales as $\Theta(nr^3(n))$.*

5 Interference-Limited Capacity

5.1 General Results

Interference can severely limit the network capacity. In this section we obtain scaling laws for the SPR, MPR and MPTR interference models. We primarily focus on deducing lower bounds. Moreover, to facilitate a succinct analysis, we develop some additional terminology and establish some important results.

Recall from Section 3 that, for two different communication schemes A and B (e.g., SPR and MPR) defined in the same graph, their corresponding interference functions I_A and I_B are different.

Definition 5.1 Dual-Interference-Set:

Consider an edge set E and an interference set $I(e)$ for an edge $e \in E$, as defined in Definition 3.3. The dual interference-set for e is defined by $F(e) = \{\hat{e} \in E \mid e \in I(\hat{e})\}$, which is the set of edges that experience a collision on account of a transmission on edge e .

Definition 5.2 Dual Conflict Graph:

Given a graph $G(V, E)$ and an interference function I , we defined the dual conflict graph as $G_D(E, E_D)$, where $E_D = \{(e, \hat{e}) \mid \hat{e} \in I(e)\}$.

Definition 5.3 Total Degree in Dual Conflict Graph:

The total degree of each node in a dual conflict graph is equal to $|M(e)|$ where $M(e) = I(e) \cup F(e)$.

Similar to the work in [16], we have the following Lemma.

Lemma 5.4 *Consider a graph $G(V, E)$ and interference I . Let $\kappa = \max_{e \in E} |M(e)|$. If f is a feasible flow rate in the absence of interference, then flow rate $f_I = f/(1 + \kappa)$ is feasible in presence of interference I .*

Proof: In the absence of interference the capacity of each edge is assumed to be one. However, because of interference, all edges cannot be activated simultaneously. Let σ

be the minimum frequency with which each edge is activated without causing any interference conflicts. Then, for each edge we have $c(e) \geq 1/\sigma$.

Observe that κ is the maximum vertex degree of the dual conflict graph G_D . It is well known that, if κ is the maximum vertex degree, then $\kappa + 1$ colors are sufficient to provide a proper vertex coloring [2]. Thus, we can partition the edge-set E into $1 + \kappa$ subsets such that no two edges in the same subset interfere. Consequently, we can periodically activate these subsets to realize $c(e) \geq 1/(1 + \kappa)$ for each edge. Thus, a feasible flow rate $f_I = f/(1 + \kappa)$ can be obtained by scaling the flow functions associated with f by a factor of $1/(1 + \kappa)$. ■

The maximum vertex degree does not provide a tight bound on the minimum number of colors required to provide a proper vertex coloring. Hence, in order to analyze a wider variety of protocol models, we introduce the concept of *interference clones*.

Definition 5.5 Interference Clone:

Two edges e_1, e_2 are said to be *interference-clones* under function I if they satisfy the conditions that $M(e_1) = M(e_2)$.

Lemma 5.6 Clone Piggy-backing Lemma: Consider a graph $G(V, E)$ along with interference functions I_A and I_B , then I_A and I_B are such that:

1. $\kappa = \max_{e \in E} |M_A(e)|$
2. $\forall e \in E$ there exists a set $M_{A, \bar{B}}(e) \subseteq M_A(e)$ such that every edge belonging $M_{A, \bar{B}}(e)$ is an *interference-clone* of e under I_B . Further, let $\mu = \min_{e \in E} |M_{A, \bar{B}}(e)|$.

If f is a feasible flow rate in G without any interference, $f_{I_A} = f/(1 + \kappa)$ is a feasible flow rate in G under the I_A interference function and κ as its corresponding parameter, then $f_{I_B} = f(1 + \mu)/(1 + \kappa)$ is feasible in presence of interference defined by I_B .

Proof: Consider the interference defined by I_A . From Lemma 5.4 we know that there exists a conflict free periodic schedule which can activate each edge at least once every $(1 + \kappa)$ slots. Let us represent this schedule by an indicator function $\alpha(e, \tau)$ which equals one iff edge e is active in slot τ . Note that the capacity of each edge under schedule α is given by

$$c_\alpha(e) = \sum_\tau \alpha(e, \tau) = 1/(1 + \kappa). \quad (38)$$

Now let us use this schedule in the presence of interference I_B . Observe that, for every e , α allocates a distinct slot for each edge in $M_{A, \bar{B}}(e) \cap \{e\}$. Consequently, every edge has $|M_{A, \bar{B}}(e)|$ interference clones scheduled in slots distinct from each other and the edge itself. In addition, note that if an edge is activated in a time slot meant for one of

its interference clones, then it does not lead to any conflict. Therefore, we can define a new conflict-free schedule β such that $\beta(e, \tau) = 1$ iff there exists an $e_1 \in M_{A, \bar{B}}(e) \cap \{e\}$ such that $\alpha(e_1, \tau) = 1$. Given that $\mu = \min_{e \in E} |M_{A, \bar{B}}(e)|$, the capacity of each edge under schedule β for interference I_B , is given by

$$\begin{aligned} c_\beta(e) &= \sum_{e_1 \in M_{A, \bar{B}}(e)} \sum_\tau \alpha(e_1, \tau) \\ &\leq (1 + \mu) \times (1/(1 + \kappa)) \end{aligned} \quad (39)$$

Accordingly, a feasible flow of $f_{I_B} = f(1 + \mu)/(1 + \kappa)$ can be obtained by scaling all the flow functions associated with the interference-free flow by a factor of $\frac{(1 + \mu)}{(1 + \kappa)}$. ■

In the subsequent discussion, we find it convenient to deduce a bound for a particular interference model and then show that it applies to a wider set of models. In order to facilitate such arguments, we define the following partial order.

Definition 5.7 Partial Order of Interference Models:

An interference function I_A is said to be *more restrictive* than I_B , represented as $I_A \preceq I_B$, iff every edge satisfies the conditions that $I_B(e) \subseteq I_A(e)$.

Lemma 5.8 Consider a graph $G(V, E)$ along with interference I_A and I_B . If $I_A \preceq I_B$, then a feasible flow rate under I_A remains feasible under I_B .

Proof: A conflict free schedule under I_A remains conflict free under I_B . Hence we can say that

$$c_A(e) \leq c_B(e) \quad (40)$$

where $c_A(e)$ and $c_B(e)$ represent the edge capacities under each interference model. Therefore, if a particular flow satisfies the capacity constraints under I_A , it necessarily satisfies those same constraints under I_B . ■

5.2 Lower Bounds

Direct analysis of the interference models can be tedious. Hence, for mathematical convenience, we consider the following restrictive models. Lemma 5.8 allows us to utilize the performance limits under these restrictive models to indirectly bind the capacity under the interference models for SPR, MPR and MPTR.

We define a restrictive interference model that introduces more restrictions (i.e., collisions) on the interference set for each edge than those strictly dictated by the original interference model. We will show that, under this restrictive model, the order of the lower bound capacity achieves the upper bound under the original (non-restrictive) interference model.

Definition 5.9 Restricted SPR (RSPR) Model:

$$I_{RSPR}(e) = W(e) - \{e\} \quad \forall e \in E_{r,H}$$

where $W(e) = \bigcup_{\hat{e}: \zeta(\hat{e}^-) = \zeta(e^-)} M_{SPR}(\hat{e})$ (41)

Definition 5.10 Restricted MPR (RMPR) Model:

$$I_{RMPR}(e) = W(e) - U(e) \quad \forall e \in E_{r,H}$$

where $U(e) = \{\hat{e} \in E_{r,H} \mid \hat{e}^- = e^-\}$ (42)

Definition 5.11 Restricted MPTR (RMPTR) Model:

$$I_{RMPTR}(e) = W(e) - V(e) \quad \forall e \in E_{r,H}$$

where $V(e) = \bigcup_{\hat{e}: \zeta(\hat{e}^-) = \zeta(e^-)} U(\hat{e})$ (43)

Consider the following properties of the restricted models.

Lemma 5.12 For the graph H_r , we have a partial order defined by (i) $I_{SPR} \preceq I_{MPR} \preceq I_{MPTR}$, (ii) $I_{RSPR} \preceq I_{RMPR} \preceq I_{RMPTR}$, (iii) $I_{SPR} \preceq I_{RSPR}$ (iv) $I_{RMPR} \preceq I_{MPR}$ (v) $I_{RMPTR} \preceq I_{MPTR}$

Proof: (Sketch) The proof for the partial orders (i) to (iii) follows directly from the definitions. Partial order (iv) follows from the fact that $J(e) \subseteq W(e)$ and $U(e) \subseteq A(e) - B(e)$. Similarly, the fact that $A(e) \subseteq V(e)$ implies partial order (v). ■

Lemma 5.13 If $r(n) \geq r_c(n)$, then all edges $e \in E_{r,H}$ have $|M(e)| = \Theta(n^2 r^4(n))$ under interference described by either of the models: $I_{SPR}, I_{MPR}, I_{MPTR}, I_{RSPR}, I_{RMPR}, I_{RMPTR}$.

Proof: From Lemma 5.12 we can conclude that RSPR is the most restrictive model, while MPTR is the least restrictive model. Further note that $F_{RSPR}(e) \subseteq W(e)$ and $I_{MPTR}(e) \subseteq M_{MPTR}(e)$. Hence, it suffices to prove the following:

$$\gamma_{max} = \max_{e \in E_{r,H}} |W(e)| = O(n^2 r^4(n)) \quad (44)$$

$$\gamma_{min} = \min_{e \in E_{r,H}} |I_{MPTR}(e)| = \Omega(n^2 r^4(n)) \quad (45)$$

Recall that a node in H_r is connected to all and only those nodes that are placed in adjacent squarelets. Hence Lemma 4.10 tells us that the degree of each vertex $v \in H_r$ is bounded as

$$4c_1 n r^2(n) \leq \deg(v) \leq 4c_2 n r^2(n) \quad (46)$$

Now lets prove the lower bound by considering the model MPTR. According to Definition 3.6, the transmission on edge e experiences interference from any transmission by

a node v such that $r(n) < \|X_{e^-} - X_v\| \leq (1 + \eta)r(n)$. Therefore, there exists an annular ring around e^- of width $\eta r(n)$ such that any transmission from a node in this ring interferes with e . The area of this annular ring is given by

$$\text{area of annular ring} = \eta(2 + \eta)\pi r^2(n) \quad (47)$$

We have already seen (Lemma 4.2) that an area of $\Theta(r^2(n))$ contains at least $\Theta(nr^2(n))$ nodes. Hence, there exists a c_3 such that

$$\gamma_{min} \geq c_3 n r^2(n) \times 4c_1 n r^2(n) \quad (48)$$

Eq. (48) proves the reqd. lower bound. The proof for the upper bound is obtained with a similar argument. First, let us inspect the transmission along edge e under the SPR model. Any transmission from a node in a disk of radius $(1 + \eta)r(n)$ around e^- , interferes with e . Moreover, a transmission from e^+ interferes with any reception in a disk of radius $(1 + \eta)r(n)$ around e^+ . Given that $\|X_{e^+} - X_{e^-}\| \leq r(n)$, there exists a disk of radius $(2 + \eta)r(n)$ around e^- containing all the nodes that may have transmission conflicts with e . Further note that a squarelet can be completely inscribed in a circum-circle of radius $(1/3\sqrt{2})r(n)$. Hence, there exists a circle of maximum radius $R = (2 + \eta + (1/3\sqrt{2}))r(n)$ containing all the nodes that may have transmission conflicts with e under the model RSPR. The area of R is $\Theta(r^2(n))$ and hence there exists a constant c_4 such that

$$\begin{aligned} \gamma_{max} &= \max_{e \in E_{r,H}} |W(e)| \\ &\leq (\text{max. no. of nodes in } R) \times (\text{max. node degree}) \\ &\leq c_4 n (2 + \eta + (\sqrt{2}/6))^2 r(n)^2 \times 4c_2 n r^2(n) \\ &= O(n^2 r^4(n)) \end{aligned} \quad (49)$$

Lemma 5.14 Consider the graph H_r with $r(n) \geq r_c(n)$ and $k \geq \Theta(n)$. In such a graph, each edge e has at least $\Theta(nr^2(n))$ clones under interference I_{RMPR} and $\Theta(n^2 r^4(n))$ clones under interference I_{RMPTR} , such that these clones interfere with each other and e , under the interference I_{RSPR} .

Proof: According to Definitions 5.10 and 5.11, $U(e)$ and $V(e)$ represent the desired set of clones for I_{RMPR} and I_{RMPTR} respectively. Lemma 4.10 implies that there exist c_1 and c_2 such that

$$\begin{aligned} \mu_{RMPR} &= \min_{e \in E_{r,H}} |U(e)| \\ &\geq \text{min. vertex degree} = c_1 n r^2(n) \end{aligned} \quad (50)$$

$$\begin{aligned} \mu_{RMPTR} &= \min_{e \in E_{r,H}} |V(e)| \\ &\geq [\min_{e \in E_{r,H}} |U(e)|] \times \text{min. nodes per squarelet} \\ &\geq c_1 n r^2(n) \times c_2 n r^2(n) \end{aligned} \quad (51)$$

■

Theorem 5.15 For $r(n) \geq r_c(n)$ and $k \geq \Theta(n)$, the capacity of random geometric network is at least: (a) $\Theta(1/r(n)k)$ under the SPR model, (b) $\Theta(nr(n)/k)$ under the MPR model, and (c) $\Theta(n^2r^3(n)/k)$ under the MPTR model.

Proof: Recall that the capacity of the random network is greater than the feasible flow rate in H_r . Theorem 4.12 shows that a rate of $f = c_1n^2r^3(n)/k$ is feasible in H_r . Additionally, Lemma 5.13 shows that the size of the largest interference set under RSPR is at most $\kappa = c_2n^2r^4(n)$. Hence, Lemma 5.4 implies that a rate of

$$f_{RSPR} = f \times (1/(1 + \kappa)) \geq (c_3/r(n)k) \quad (52)$$

is feasible under the RSPR model. If we take into consideration the interference clones, then Lemma 5.6 further implies that the rate

$$\begin{aligned} f_{RMPR} &= f \times (1/(1 + \kappa)) \times (1 + \mu_{RMPR}) \\ &\geq (c_3/r(n)k) \times (c_4nr^2(n)) \\ &= (c_3c_4nr(n)/k) \end{aligned} \quad (53)$$

is feasible under the RMPR model. Similarly, the rate

$$\begin{aligned} f_{RMPTR} &= f \times (1/(1 + \kappa)) \times (1 + \mu_{RMPTR}) \\ &\geq (c_3/r(n)k) \times (c_5n^2r^4(n)) \\ &= (c_3c_5n^2r^3(n)/k) \end{aligned} \quad (54)$$

is feasible under the RMPTR model. Finally, note that a feasible rate under a restricted model is necessarily feasible under the original model. Hence, the result proven in Lemma 5.13 completes the proof. ■

5.3 Upper Bounds

The interference-free capacity provides an upper bound on the capacity under any model, and Theorem 5.15 already shows that the MPTR model achieves this capacity. However, we need to provide additional arguments to obtain a tight bound on the capacity under SPR and MPR. Our arguments are similar to [15] and [6], and we briefly sketch the proof of the following result for the sake of completeness.

Theorem 5.16 For $r(n) \geq r_c(n)$ and $k \geq \Theta(n)$, the capacity of random geometric network is at most: (a) $\Theta(1/r(n)k)$ under the SPR model, (b) $\Theta(nr(n)/k)$ under the MPR model, and (c) $\Theta(n^2r^3(n)/k)$ under the MPTR model.

Proof: (Sketch) Consider the Cut S in Figure 1. The capacity of this cut and hence the network is less than

$$\frac{(\text{no. of transmitters in } A) \times (\text{max. transmission per node})}{(1/2)k} \quad (55)$$

The total number of transmitters in A is less than the total number of nodes in A . We have already shown that the total number of nodes is $\Theta(nr(n))$. Under the MPTR model, each node can transmit a maximum of $\Theta(nr^2(n))$ packets, and under the MPR model each packet transmits just a single packet. This provides the bounds for MPR and MPTR. To obtain the bound for SPR, we note that each transmitter silences all the nodes within an area of $\Theta(r^2(n))$. Hence, only $\Theta(1/r(n))$ nodes are capable of transmitting simultaneously across a cut. Each of these nodes transmits a single packet and, therefore, the above equation provides the bound for SPR. ■

6 Discussion

The results we have presented demonstrate that future ad hoc networks can scale well beyond the Gupta-Kumar capacity bounds attained when nodes simply try to avoid MAI [8]. First, we showed that the optimal capacity that any protocol architecture can attain in a wireless network is $\Theta(n^2r^3(n)/k)$. Second, we demonstrated that this capacity can indeed be attained when nodes embrace MAI as transmitters and receivers, and that a non-vanishing capacity is attainable per S-D pair even when information must be disseminated over multiple hops. Given that the majority of nodes in an ad hoc network are sources and destinations of information, we frame the rest of our discussion for the important case of $k = \Theta(n)$.

The capacity under the SPR model when $k = \Theta(n)$ equals $\Theta(1/nr(n))$, which is the well-known Gupta-Kumar result, and shows that avoiding MAI in the communication protocols of ad hoc networks does not scale. Given that today's ad hoc networks are based on interference avoidance, it is clear that massively scalable ad hoc networks cannot exist without substantial changes to the way in which protocols are designed.




Figure 4. Numerical results

On the other hand, under the MPR model [6] the network capacity is equal to $\Theta(r(n))$, which constitutes a dramatic improvement over SPR, and promises to provide capacity

gains in practice by embracing interference at the receivers. However, to attain a non-vanishing capacity, $r(n)$ must be $\Theta(1)$, i.e., use single-hop communication. Unfortunately, this is not feasible in practice, because of the energy that would be incurred in transmissions and the complexity required for the receivers to decode a number of transmissions in the order of the nodes in the network.

In contrast to the above, MPTR achieves the optimal per-node capacity of $\Theta(nr^3(n))$. Thus, any choice of $r(n) = \Omega(n^{-1/3})$ allows us to increase the per-node capacity of the network with n , while still having multihop communication. Moreover, the transmission range and hop-size decreases with n . As per our analysis, the capacity of a network is a constant factor of Γ , where $\Gamma = 2\eta_{min}/l(k_{max})$, such that η_{min} is the minimum number of edges between any two squarelets, k_{max} is the maximum number of sources or destinations in a single squarelet and $l \times l$ is the total number of squarelets. To illustrate this, we numerically evaluated the behavior of Γ as a function of n . Figure 4 presents the mean, minimum and maximum observed value of Γ over a thousand network topologies randomly generated and in which $k = n/2$ and $r(n) = 1/n^{0.25}$. Clearly, as $n \rightarrow \infty$ we have $r(n) \rightarrow 0$. Nevertheless, the numerical results show that the per-node capacity still increases as $\Theta(n^{0.25})$.

In closing, we should point out that, while our results provide a completely new outlook on the design of wireless ad hoc networks, much work remains to be done to fully understand their fundamental limits! For example, the results we have presented address only unicast traffic; our model can be used to study the cases of multicast and broadcast information dissemination. We also hope that this paper motivates research on protocol architectures that combine multi-packet reception and transmission to attain massively scalable ad hoc networks.

7 Acknowledgments

This work was partially sponsored by the U.S. Army Research Office under grants W911NF-04-1-0224 and W911NF-05-1-0246, by the National Science Foundation under grant CCF-0729230, by the Defense Advanced Research Projects Agency through Air Force Research Laboratory Contract FA8750-07-C-0169, and by the Baskin Chair of Computer Engineering. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

References

[1] R. Ahlswede, C. Ning, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, 2000.
 [2] B. Bollobas. *Modern Graph Theory*. Springer Verlag, 1998.

[3] R. M. de Moraes, H. R. Sadjadpour, and J. J. Garcia-Luna-Aceves. Many-to-many communication: A new approach for collaboration in manets. In *Proc. of IEEE INFOCOM 2007*, Anchorage, Alaska, USA., May 6-12 2007.
 [4] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton Univ. Press, 1962.
 [5] M. Franceschetti, O. Dousse, D. Tse, and P. Thiran. Closing the gap in the capacity of wireless networks via percolation theory. *IEEE Transactions on Information Theory*, 53(3):1009–1018, 2007.
 [6] J. J. Garcia-Luna-Aceves, H. R. Sadjadpour, and Z. Wang. Challenges: Towards truly scalable ad hoc networks. In *Proc. of ACM MobiCom 2007*, Montreal, Quebec, Canada, September 9-14 2007.
 [7] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, 2002.
 [8] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.
 [9] S. Katti, S. Gollakota, and D. Katabi. Embracing wireless interference: Analog network coding. In *Proc. of ACM SIGCOMM 2007*, Kyoto, Japan, August 27-31 2007.
 [10] P. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proc. of ACM symposium on Theory of computing*, San Diego, California, USA, May 16-18 1993.
 [11] S. Kulkarni and P. Viswanath. A deterministic approach to throughput scaling wireless networks. *IEEE Transactions on Information Theory*, 50(6):1041–1049, 2004.
 [12] M. Kunde. Block gossiping on grids and tori: Deterministic sorting and routing match the bisection bound. In *Proc. of European Symp. Algorithms*, 1991.
 [13] P. Kyasanur and N. Vaidya. Capacity of multi-channel wireless networks: Impact of number of channels and interfaces. In *Proc. of ACM MobiCom 2005*, Cologne, Germany, August 28-September 2 2005.
 [14] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999.
 [15] J. Liu, D. Goeckel, and D. Towsley. Bounds on the gain of network coding and broadcasting in wireless networks. In *Proc. of IEEE INFOCOM 2007*, Anchorage, Alaska, USA., May 6-12 2007.
 [16] R. Madan, D. Shah, and O. Leveque. Product multicommodity flow in wireless networks. *Submitted to IEEE Transactions on Information Theory*, 2007.
 [17] A. Ozgur, O. Leveque, and D. Tse. Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks. *IEEE Transactions on Information Theory*, 53(10):2549–3572, 2007.
 [18] S. Toumpis and A. J. Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE Transactions on Wireless Communications*, 2(4):736–748, 2003.
 [19] Y. Wu, P. A. Chou, and S.-Y. Kung. Information exchange in wireless networks with network coding and physical-layer broadcast. In *Proc. of CISS 2005*, Baltimore, MD, March 16-18 2005.