The User Capacity of Barrage Relay Networks

Thomas R. Halford
TrellisWare Technologies, Inc.
16516 Via Esprillo, Suite 300
San Diego, CA 92127
Email: thalford@trellisware.com

Thomas A. Courtade
Center for Science of Information
Stanford University
Stanford, CA 94305
Email: tom.courtade@gmail.com

Kurt A. Turck
Networking Technology Branch
Air Force Research Laboratory
Rome, NY 13441
Email: Kurt.Turck@rl.af.mil

***Abstract*—Barrage relay networks (BRNs) are a class of mobile ad hoc networks based on an autonomous cooperative communications scheme that affords a distributed, rapid, and robust broadcast mechanism. BRN-based radios are currently being used operationally; understanding scaling laws for BRNs can thus shed light on how future systems ought be designed to address a wider range of military missions. It has previously been shown that BRNs scale optimally for broadcast traffic (in terms of sum throughput and latency). Furthermore, experimental evidence supports the scalability of BRNs in practice: a 215-node network of TrellisWare's BRN-based CheetahNet radios was demonstrated in 2010 as part of the Rim of the Pacific (RIMPAC) exercise.**

In this work, we study scaling laws for unicast in BRNs. The spatial extent of unicast transmissions in BRNs can be contained via controlled barrage regions (CBRs), while barrage access control protocols – which are a type of path-oriented medium access control (PO-MAC) protocol – can be used to schedule CBRs in space and time. Traditional techniques for scaling analysis support neither autonomous cooperation nor PO-MAC protocols; we therefore study the fundamental limits of CBR-based protocols through the lens of *user capacity*, which we define as the maximum number of constant-rate, delay-optimal unicast flows that an n-node network can simultaneously support.

The user capacity of dense networks operating under a CBR protocol is shown to scale as $\Theta(\sqrt{n})$ when source-destination pairs are chosen randomly. BRNs are thus able to transport $\Theta(\sqrt{n/\log n})$ bit-meters/sec, which is order optimal (in the Gupta-Kumar sense). If the source-destination pairs are instead chosen so that their separation (in hops) is a geometric random variable, then the user capacity scales as $\Theta(n/\log n)$. Thus, in a wireless network dominated by localized traffic, BRNs can achieve throughput that is almost linear in the number of users.

Finally, it is shown via software simulation that distributed, greedy algorithms provide order optimal scheduling of CBRs for both the random and localized traffic models. This suggests that scalable protocols for unicast in BRNs can be realized in practice.

I. Introduction

In response to the limited success of traditional, strictly layered approaches to mobile ad hoc network (MANET) architectures for military applications (cf., [1–3]), barrage relay networks were designed from the ground up to meet the demands of communications at the tactical edge [4]. As detailed in [5], BRNs utilize time division multiple access (TDMA) and cooperative communications as the basis of an efficient broadcast protocol wherein packets ripple out from a source in pipelined spatial waves. This broadcast mechanism is the fundamental physical layer resource in a BRN; barrage access control (BAC) protocols operating at Layer 2 coordinate which node is the broadcast source on any given time slot, rather than which point-to-point links are active. It was shown in [5] that BRNs provide optimal broadcast scaling: in a network where each node can transmit W bits/sec, the sum throughput that can be achieved by an arbitrary number of broadcast sources scales as[1] $\Theta(W)$.

BRNs were proposed initially for platoon-sized missions dominated by latency-critical, broadcast and multicast traffic (e.g., push-to-talk (PTT) voice and real-time streaming video). Increasingly, BRN-based radios are being used to support larger missions and richer classes of data services. For example, the Marine Corps Warfighting Laboratory deployed a 215-node network of TrellisWare's CheetahNet radios as part of the Limited Objective Experiment 4 (LOE-4) exercise at RIMPAC 2010. During this exercise, applications that are inherently unicast (e.g., IP chat) were supported in addition to broadcast services such as PTT voice and position location information (PLI) monitoring. Such practical use cases motivate the study of efficient protocols for unicast transport in BRNs.

The building blocks for unicast protocol design in BRNs have been described in [5–8]. Briefly, the spatial extent of transmissions in BRNs can be be contained via controlled barrage regions (CBRs). BAC protocols can be used to schedule CBRs that are multiplexed in both space and time (i.e., more than one unicast source is active on any time slot). This work studies the fundamental limits that govern *any* protocol employing CBRs for unicast transmission in BRNs.

The assumptions made in traditional scaling studies [9] fail to capture two salient features of BRNs: cooperative communications and path-oriented medium access control techniques that reserve entire multihop routes at Layer 2. We therefore study scaling laws for BRNs through the lens of *user capacity*, which we define as the maximum number of constant-rate, delay-optimal unicast flows that an n-node

Approved for Public Release; Distribution Unlimited: 88ABW-2012-4037 dated 20 Jul 2012. ACKNOWLEDGMENT OF SUPPORT AND DISCLAIMER: (a) Contractor acknowledges Government's support in the publication of this paper. This material is based upon work funded by AFRL under AFRL Contract No. FA8750-11-C-0169. (b) Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of AFRL.

[1]A function $f(n) = \Omega(g(n))$ if $\exists$ constants n_0 and c such that $f(n) \geq cg(n), \forall\, n > n_0$. Similarly, $f(n) = O(g(n))$ if $f(n) \leq cg(n), \forall\, n > n_0$. If $f(n)$ is $\Omega(g(n))$ and $O(g(n))$, then we say $f(n) = \Theta(g(n))$. Throughout, $o(1)$ denotes a quantity that vanishes as $n \to \infty$.

network can simultaneously support. In the context of BAC protocols, user capacity can be interpreted as the maximum number of CBRs that can be accessed at any given time. In a tactical scenario, user capacity measures, for example, how many streaming video feeds can be simultaneously supported, a measure not readily derived from the analysis of [9].

Using several new results that are derived in an appendix, we demonstrate in Section III that the user capacity of BRNs employing CBRs scales as $\Theta(\sqrt{n})$ when the source and destination nodes are paired randomly. BRNs are thus capable of transporting $\Theta(\sqrt{n/\log n})$ bit-meters/sec and can achieve the transport capacity defined by Gupta and Kumar [9]. Intuitively, this result indicates that although a given CBR contains more nodes than is strictly required for unicast transport – the CBR contains all possible shortest paths from the source to the destination, rather than a single shortest path – this does not affect asymptotic scalability (all the while dramatically increasing the reliability [10] and stability [8] of multihop data transport). It is further demonstrated in Section III that the user capacity of BRNs employing CBRs scales as $\Theta(n/\log n)$ when the distances between unicast source/destination pairs is geometrically distributed. Thus, in a wireless network dominated by localized traffic, BRNs can achieve throughput that is almost linear in the number of users.

Finally, it is shown via software simulation in Section III that distributed, greedy algorithms provide order optimal scheduling of CBRs for both the random and localized traffic models. This suggests that scalable protocols for unicast transport in BRNs can be realized in practice. This paper concludes in Section IV with a discussion of the design challenges associated with such practical protocols.

II. Models and Metrics

A. Network Model

We adopt the standard random network model in this work wherein n nodes are distributed randomly on the unit square and communication at a rate of W bits/sec is possible between two nodes if and only if they are separated by less than the transmission range $r(n)$. A key result in geometric random graph theory asserts that the threshold on $r(n)$ required for such a network to be connected scales as [11, 12]:

$$r(n) \sim \Theta\left(\sqrt{\frac{\log n}{n}}\right). \tag{1}$$

This critical density is assumed throughout. When simulating random networks in Section III, a transmission range of

$$r^\star(n) = \frac{3}{2}\sqrt{\frac{\log n}{\pi n}} \tag{2}$$

is assumed to assure connectivity with high probability[2].

[2]An event E occurs with high probability if $\forall\, \alpha > 1$ there is an appropriate choice of constants for which E occurs with probability at least $1-O(1/n^\alpha)$.

B. Controlled Barrage Regions and CBR Protocols

Due to space limitations, the reader is referred to [5] for a full description of BRNs. It suffices presently to describe CBRs. Briefly, controlled barrage regions can be established by a reactive protocol as follows. Control messages broadcast by the source (s) and destination (d) of a unicast flow inform nodes of their respective distance to s and d in hops, as well as the distance of the shortest path from s to d. Nodes then use these three distances to ascertain whether they are on *any* of the shortest paths from s to d, in which case they are interior to the CBR and relay messages from s per the barrage cooperative decode-and-forward protocol. Nodes that are not interior to the CBR, but which are neighbors of interior nodes, suppress their relay function so that external packets do not propagate into the CBR, nor do internal packets propagate to the rest of the network. In this way, multiple unicast transmissions may be established in the network.

Definition 1 (CBR Protocol): A *CBR protocol* is a mechanism that establishes a controlled barrage region such that a node is interior to the CBR if and only if it lies on *any* shortest path connecting the source and destination.

C. Path-Oriented Medium Access Control

Although we focus primarily on BRNs, the results presented herein can be interpreted in the context of *any* PO-MAC protocol that reserves entire source-destination (s-d) paths at Layer 2 (e.g., [1]). Observe that an instantaneous throughput of $W/3$ bits/sec can be achieved along what is effectively a multihop line network under such protocols. Furthermore, this transmission is delay-optimal: the end-to-end latency depends only on the average s-d hop distance and not on, for example, the number of two-hop neighbors of any relay node.

Our goal in this paper is to compare the CBR protocol to the performance attainable by any network that employs a PO-MAC protocol. To this end, we make the following definition:

Definition 2 (PO-MAC User Capacity): The *PO-MAC user capacity* is the maximum number of mutually non-interfering unicast flows that can be simultaneously supported by a wireless network using a PO-MAC protocol. Observe that these flows support $\Theta(W)$-throughput, delay optimal data transport provided the PO-MAC reserves entire routes.

We show in the Section III that CBR protocols attain order optimal scaling with respect to the PO-MAC user capacity.

D. Unicast Traffic Models

Finally, we consider two different types of traffic model which represent the extremes of "global" and "local" traffic in this paper. In a realistic scenario, the traffic generation process would likely be a mixture of these two extremes.

Definition 3 (Uniform Traffic): Unicast traffic is said to be generated according to the *uniform traffic model* if $n/2$ distinct source-destination pairs are chosen randomly from the n network nodes.

Note that the uniform traffic model is equivalent to randomly placing $n/2$ source-destination pairs in the plane and generating the induced random geometric graph.

Definition 4 (Localized Traffic): Unicast traffic is said to be generated according to the *localized traffic model with mean* β if $n/2$ distinct sources are chosen from the set of vertices, and for each source s, a destination is chosen uniformly from the set of vertices that are $H(u)$ hops away from s. Here, $H(u)$ is a geometric random variable with mean β.

In the localized traffic model, it is possible that a node is a source for one unicast flow and a destination for another distinct – albeit mutually interfering – flow.

III. Main Results and Discussion

A. Results

All of our results assume that the underlying network is modeled as a dense random geometric graph with connectivity radius $r(n) = \Theta(\sqrt{\log n/n})$ chosen sufficiently large to ensure connectivity with high probability (w.h.p.). All results hold with high probability in n. With the exception of the first result, we delay proofs until the appendix.

Main Result 1: Under any traffic model, a network operating under a PO-MAC protocol supports at most $O(n/\log n)$ simultaneous unicast flows.

Proof: A PO-MAC protocol reserves paths in a network, one per unicast flow, and the paths must be mutually non-interfering. Each path creates an interference region of area greater than $\pi r(n)^2$ since it must contain at least one transmitting node. Dividing the area of the unit square by $\pi r(n)^2$ yields the desired upper bound. ∎

This first main result places an upper limit on the user capacity of any network employing a PO-MAC protocol, and hence an upper limit on the number of CBRs that can be accessed simultaneously in a BRN. Intuitively, this upper bound results from the "interference footprint" of the source nodes alone. Our second main result indicates that this upper bound is in fact achievable in BRNs when traffic is localized:

Main Result 2: A BRN network using a CBR protocol supports $\Theta(n/\log n)$ simultaneous flows under the assumptions of the random localized traffic model.

Note that Result 1 does not imply that *any* $\Theta(n/\log n)$-sized subset of the $n/2$ flows in a realization of the localized traffic model can be supported. Instead, it says that the size of the largest simultaneously supportable subset scales as $\Theta(n/\log n)$. Result 1 can be combined with the lower bound implicit in Result 2 – i.e., if a CBR protocol achieves $\Omega(n/\log n)$ user capacity then at least one path-oriented MAC protocol has been shown to do so – to yield a scaling law that characterizes the user capacity for networks operating with *any* PO-MAC protocol under the localized traffic model.

Main Result 3: Under the localized traffic model, the PO-MAC user capacity scales as $\Theta(n/\log n)$.

Results 2 and 3 together are interesting as they demonstrate that CBR protocols scale as well as any path-oriented MAC protocol. That is to say, the additional relay nodes reserved in a CBR beyond the shortest path route nodes do not affect asymptotic scalability when traffic is localized. It was shown in [10] and [8] that these additional relay nodes serve to enhance the reliability and stability, respectively, of multihop data transport, even over 2- and 3-hop CBRs.

Our last two analytical results concern the transport capacity of BRNs and the number of unicast flows that can be supported by a BRN under the assumptions of the uniform traffic model.

Main Result 4: Under the uniform traffic model, a BRN employing a CBR protocol supports $\Theta(\sqrt{n})$ simultaneous unicast flows.

Main Result 5: Under the uniform traffic model, a BRN employing a CBR protocol is capable of transporting[3] $\Theta(\sqrt{n/\log n})$ bit-meters/sec and therefore can achieve the transport capacity in the sense of [9].

Result 5 shows that, up to some constant factor that is independent of the network size n, BRNs can transport as much data as arbitrary wireless networks architectures.

Since shortest paths between vertices in a sufficiently dense random geometric graph can be approximated by straight-line segments, the results derived in the next section suggest that under the uniform traffic model, a network operating under a PO-MAC protocol can support at most $O(\sqrt{n})$ simultaneous traffic flows. This leads to the following conjecture:

Conjecture 1: A barrage relay network employing a CBR protocol achieves the PO-MAC user capacity (in the sense of order-optimal scaling) for the random uniform traffic model.

The proofs of these results hint that a greedy CBR formation protocol can be employed to achieve the user capacity. To this end, we have simulated a greedy CBR formation protocol under each traffic model. The numerical results strongly indicate that the performance of a greedy CBR formation protocol coincides with the optimal scaling results given in the above results. Simulation results are given in Figure 1.

B. Discussion

By their very nature, BRNs embrace the fundamental broadcast characteristic of the the wireless medium rather than try to control it via a collision avoidance mechanism. It was demonstrated in [5] that a simple Layer 2 protocol achieves order optimal scaling for broadcast traffic. However, if the rapid broadcast mechanism is left unchecked, BRNs suffer from scalability issues for unicast and multicast traffic since a single multicast intended for a small group of receivers can monopolize the entire network for the duration of the transmission. Employing a protocol that uses CBRs to contain the spatial extent of BRN broadcasts is one method of coping with this scalability issue. While the CBR protocol is a natural extension of the traditional barrage framework, it was unknown until now whether it would scale in an optimal manner for unicast traffic. The results of the previous subsection indicate that the answer to this question is affirmative for the localized traffic model, and numerical and analytical evidence strongly indicate that this is the case for the uniform traffic

[3]A *bit-meter* corresponds to the transport of 1 bit of information 1 meter closer to its intended destination.

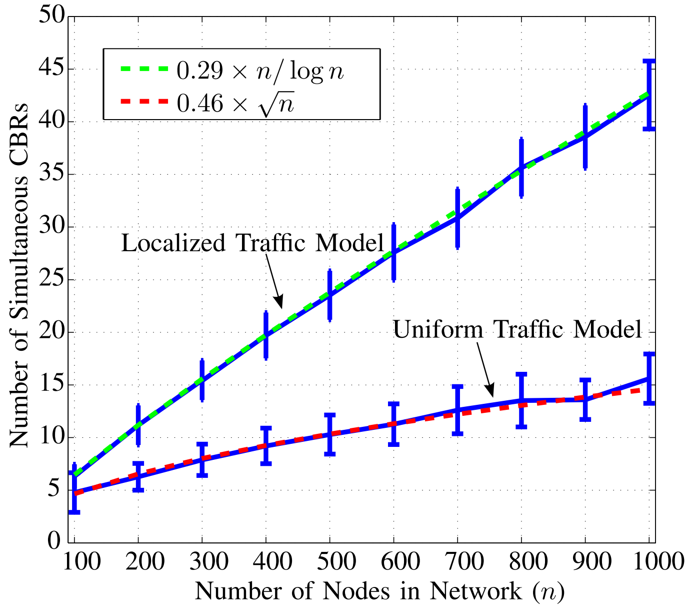

Fig. 1. Simulation results for a distributed, greedy implementation of the CBR protocol. The solid blue lines represent the number of simultaneous mutually non-interfering CBRs supported by a greedy implementation of the CBR protocol. The results were averaged over 100 independent trials for each network size n. The vertical blue lines indicate ± 1 standard deviation.

model as well. Thus, BRNs employing a CBR protocol enjoy the simplicity and benefits of a traditional BRN (cf., [5]), but they also support scaling to hundreds or thousands of nodes.

IV. Conclusion

In this paper we proved that BRNs employing a CBR protocol can support $\Theta(n/\log n)$ and $\Theta(\sqrt{n})$ simultaneous unicast sessions under localized and uniform traffic models, respectively. These results prove that order-optimal throughput scaling is achieved for the localized traffic model, and the numerical and analytical results strongly suggest the same statement is true for the uniform traffic model. We have also shown that the CBR protocol can achieve the transport capacity for the uniform traffic model. These results resolve the scaling issue for dense BRNs; future work will consider sparse networks from a theoretical perspective. Finally, we also demonstrate that a greedy, distributed implementation of the CBR protocol achieves the order-optimal scaling predicted by the theoretical results. This suggests that scalable barrage relay network architectures can indeed be realized in practice.

We have focused only on *data scalability* in this paper – i.e., how efficiently a network uses the RF spectrum when transmitting Layer 4 data. The Layer 2 and 3 control overhead that is neglected in such analysis can be a significant, even limiting, factor in practical networks. Indeed, the design of practical BAC protocols for barrage relay networks requires that the overhead used to multiplex CBRs in space and time be contained. While the mathematical tools to address networking scaling in terms of both data and control are as yet nascsent (cf., [3]), the development of a theoretical framework for understanding *control scalability* in wireless networks will provide valuable metrics against which practical protocol implementers can measure their designs.

V. Acknowledgements

We wish to thank Mike Neely of the University of Southern California for suggesting the proof technique for Theorem 3, as well as his conjecture that Claim 2 ought be provable.

Appendix

The main results presented in Section III can be immediately deduced from Theorems 1-3, all of which assume the dense random graph network model described in Section II.

Theorem 1: Under a localized traffic model, any PO-MAC protocol that reserves all the shortest paths between source-destination (S-D) pairs (e.g., a CBR protocol) can support $\Omega(n/\log n)$ simultaneous traffic flows w.h.p.

Proof: Let $H(u)$ be the (random) number of hops separating a source-destination pair. By Markov's inequality, $\Pr(H(u) \leq 2\beta) \geq 1/2$, where $\beta = \mathbb{E}[H(u)]$, and thus (w.h.p.) there exists $n/8$ S-D pairs that are at most 2β hops apart. Each of these S-D pairs, and all the shortest paths connecting the source and destination, can each be enclosed within a ball of radius $2\beta r$ centered at the respective source.

Now, tessellate the unit square by boxes with side length $2\sqrt{8\log n/n}$. By Lemma 5.7 in [13], each of these boxes contains at least one source vertex w.h.p. Now, partition the boxes into classes of equal size so that any two boxes in the same class are at least distance $5\beta r$ apart. Since each box has dimension proportional to r, a finite number of classes suffices (depending on β, but not n). The sources contained in a single class, the corresponding destinations, and the corresponding shortest paths are mutually non-interfering by construction. Thus, any PO-MAC protocol that reserves shortest paths between source-destination pairs can support $\Omega(n/\log n)$ simultaneous traffic flows w.h.p in n. ∎

By combining Theorem 1 with Main Result 1 of Section III, we obtain the following result:

Corollary 1: Under the localized traffic model, the PO-MAC user capacity scales as $\Theta(n/\log n)$. Moreover, any PO-MAC protocol that reserves shortest paths between S-D pairs scales order-optimally with respect to PO-MAC user capacity.

Now, we turn our attention toward the task of showing that BRNs with a CBR protocol can support $\Theta(\sqrt{n})$ simultaneous flows when the source-destination pairs are chosen randomly. Before we can accomplish this, we need two auxiliary lemmas. The first lemma is a combinatorial result that characterizes the number of non-crossing line segments in the unit square. This result is proved via the probabilistic method (cf., [14]). The second lemma relates line segments to paths in a random geometric graph. We begin with some definitions.

Let $Q = [0,1) \times [0,1)$ denote the unit square and let $\mathcal{L}_n$ be a set of n line segments in Q. Define the distance between two line segments $\ell_1, \ell_2 \in \mathcal{L}_n$ as the minimum distance between any point in ℓ_1 and any point in ℓ_2. Formally,

$$d(\ell_1, \ell_2) := \inf_{x\in\ell_1, y\in\ell_2} \|x - y\|,$$

where $\|x-y\|$ is the Euclidean distance between points $x, y \in Q$. A set $\mathcal{U}$ of line segments is said to be d-disjoint if all pairs of line segments in $\mathcal{U}$ are at least distance d apart. Then

$$N_d(\mathcal{L}_n) := \max_{\mathcal{U} \subseteq \mathcal{L}_n} \{|\mathcal{U}| : d(\ell_1, \ell_2) \geq d \text{ for all } \ell_1, \ell_2 \in \mathcal{U}\}$$

is the size of the largest d-disjoint subset of $\mathcal{L}_n$.

In what follows, we will distinguish between *left* and *right* endpoints of a line segment. This assignment is an artificial one, however, it simplifies the exposition considerably.

Definition 5: A line segment is said to be drawn according to a *uniform process* if the left and right endpoints are chosen independently and uniformly from the points in Q.

Lemma 1: If $\mathcal{L}_n$ is a set of n line segments generated according to a uniform process, then with probability approaching 1 as $n \to \infty$, $N_{d(n)}(\mathcal{L}_n) = \Theta(\sqrt{n})$ for any $d(n) = O(n^{-1/4})$.

Proof: We prove Lemma 1 by showing that $N_{d(n)}(\mathcal{L}_n) = \Omega(\sqrt{n})$ and $N_{d(n)}(\mathcal{L}_n) = O(\sqrt{n})$.

Claim 1: With probability tending to 1 as $n \to \infty$, $N_{d(n)}(\mathcal{L}_n) = \Omega(\sqrt{n})$ for any $d(n) = O(n^{-1/4})$.

Proof of Claim 1: Partition Q into $\sqrt{n}$ disjoint squares, each of size $n^{-1/4} \times n^{-1/4}$. Note that if a line segment ℓ is contained in a single square, then it will not intersect line segments contained in any other square. Let Y be the number of squares that contain line segments. By assumption, there exists an integer m such that $d(n) \leq mn^{-1/4}$ for all n. Therefore, by the pigeon-hole principle, we can find at least $Y(1-o(1))/(m+1)^2$ squares containing line segments that are $d(n)$-disjoint. Thus, it remains to be shown that $Y \geq c\sqrt{n}$ with high probability for some absolute constant c.

Since the area of each square is $1/\sqrt{n}$, the probability that an endpoint falls into a particular square j is $1/\sqrt{n}$. Since endpoints are chosen independently, the probability that a given line l falls entirely inside that square l is $1/n$, and the probability that square j does not contain any line segments is $(1-1/n)^n \approx 1/e$. Further, note that:

$$\begin{aligned}
&\Pr\left[\{\text{line } \ell \text{ not in square } i\} \wedge \{\text{line } \ell \text{ not in square } j\}\right] \\
&= 1 - \Pr\left[\{\text{line } \ell \text{ in square } i\} \vee \{\text{line } \ell \text{ in square } j\}\right] \\
&= 1 - \left(\Pr\left[\text{line } \ell \text{ in square } i\right] + \Pr\left[\text{line } \ell \text{ in square } j\right]\right) \\
&= 1 - 2/n,
\end{aligned}$$

where we used the fact that the events $\{\text{line } \ell \text{ in square } i\}$ and $\{\text{line } \ell \text{ in square } j\}$ are disjoint. Then for any pair of squares (i, j), the probability that neither square i nor square j contains any line segments is $(1-2/n)^n \approx 1/e^2$.

Let X_i be the indicator random variable taking value 1 if square i contains no line segments and 0 otherwise. Note that $\mathbb{E}X_i \approx e^{-1}$, $\text{Var}(X_i) \approx e^{-1}(1-e^{-1})$, and

$$\begin{aligned}
\text{Cov}(X_i, X_j) &= \mathbb{E}\left[X_i X_j\right] - \mathbb{E}\left[X_i\right]\mathbb{E}\left[X_j\right] \\
&= (1-2/n)^n - (1-1/n)^{2n} \\
&\leq e^{-2} - e^{-2} + o(1) = o(1).
\end{aligned}$$

Then, letting $X = \sum_{i=1}^{\sqrt{n}} X_i$ be the number of squares that don't contain any line segments, and noting that

$$\begin{aligned}
\text{Var}(X) &= \sum_i \text{Var}(X_i) + \sum_{i \neq j} \text{Cov}(X_i, X_j) \\
&\leq \sqrt{n}\left(\frac{1}{e}\left(1 - \frac{1}{e}\right) + o(1)\right) + \sqrt{n}(\sqrt{n}-1)o(1) \\
&\leq \sqrt{n} + n \cdot o(1),
\end{aligned}$$

Chebyshev's inequality yields:

$$\begin{aligned}
\Pr\left[|X - \mathbb{E}X| \geq \frac{1}{10}\sqrt{n}\right] &\leq 100 \times \frac{\text{Var}(X)}{n} \\
&\leq 100 \times \frac{\sqrt{n} + n \cdot o(1)}{n} \to 0.
\end{aligned}$$

Therefore, with probability tending to 1,

$$X \leq \left(1 + \frac{1}{10} + o(1)\right)\frac{\sqrt{n}}{e} \leq \frac{\sqrt{n}}{2}.$$

Hence, $Y = \sqrt{n} - X \geq \sqrt{n}/2$ with probability tending to 1. This proves the claim.

Claim 2: With probability tending to 1 as $n \to \infty$, $N_{d(n)}(\mathcal{L}_n) = O(\sqrt{n})$ for any $d(n) = O(n^{-1/4})$.

Proof of Claim 2: First, observe that it suffices to prove that $N_0(\mathcal{L}_n) = O(\sqrt{n})$ since $N_d(\mathcal{L}_n)$ is nonincreasing in d. To this end, from [15], there exists an absolute constant c such that for any $2k$ points in the plane, the number of non-crossing left-right[4] perfect matchings is upper-bounded by $c \cdot 29^k$. Consider any realization of n line segments in the plane and further consider the $2k$ (k left and k right) endpoints corresponding to any subset S consisting of k line segments. Every left-right perfect matching of these $2k$ points is equally likely, and thus the probability that these k segments are non-crossing is upper bounded by $c \cdot 29^k/k!$ since there are $k!$ left-right perfect matchings on the $2k$ endpoints. Upper bounding the right hand side of this expression using Stirling's formula and recalling the crude upper bound $\binom{n}{k} \leq \left(\frac{n \cdot e}{k}\right)^k$ yields the following upper bound:

$$\begin{aligned}
\Pr[\exists\ k \text{ non-crossing segments}] &\leq \binom{n}{k}\frac{c \cdot 29^k}{k!} \\
&\leq o(1) \cdot \left(\frac{29 \cdot n \cdot e^2}{k^2}\right)^k
\end{aligned}$$

Letting $k = 15\sqrt{n}$, we have

$$\begin{aligned}
\Pr[\exists\ 15\sqrt{n} \text{ non-crossing segments}] &\leq o(1) \cdot \left(\frac{29 \cdot e^2}{15^2}\right)^{15\sqrt{n}} \\
&\leq o(1) \cdot (.96)^{15\sqrt{n}} \to 0,
\end{aligned}$$

thus completing the proofs of the claim and Lemma 1. ■

[4] A matching in a graph is a set of pairwise non-adjacent edges – i.e., no two edges share a common vertex – while a *perfect* matching is one that matches every vertex in the graph. A left-right perfect matching simply distinguishes between left endpoints and right endpoints in edges. In other words, an edge is only allowed to match a left endpoint to a right endpoint.

Now, we state a result from [16] that relates line-segments to shortest paths in a random geometric graph.

Lemma 2 ([16] p. 426): For any two vertices v_1, v_2 in a dense random graph G, w.h.p.,

$$d_G(v_1, v_2) \leq \frac{K\|v_1 - v_2\|}{r(n)}$$

where $d_G(v_1, v_2)$ is the length of the shortest path connecting vertices v_1 and v_2 and $r(n)$ is the connectivity radius.

Lemma 2 essentially states that the length of a shortest path between two vertices is proportional to the Euclidean distance between the two corresponding points in the plane. Therefore, we can use this lemma to conclude that the shortest paths connecting two vertices must be "well-behaved". With the above lemmas in hand, we can now prove Main Result 4.

Theorem 2: Consider a set of $n/2$ randomly chosen distinct source-destination pairs in a dense random geometric graph G. With probability tending to 1 as $n \to \infty$, there exist $\Omega(\sqrt{n})$ S-D pairs which generate mutually non-interfering CBRs.

Proof: For a source-destination pair (s, d), denote the straight-line segment connecting s to d as $\ell_{s,d}$. Since the vertices of G are uniformly chosen from Q and the S-D pairs are chosen randomly, the set of segments connecting S-D pairs is generated according to a uniform process. Define

$$T(s, d, \alpha) = \left\{ x \in Q : \inf_{y \in \ell_{s,d}} \|x - y\| \leq \alpha \right\}$$

to be the set of points in Q that are within distance α of $\ell_{s,d}$. Lemma 1 states that, with high probability, there exist $\Theta(\sqrt{n})$ distinct source-destination pairs such that $T(s, d, n^{-1/4}) \cap T(s', d', n^{-1/4}) = \emptyset$ for distinct pairs (s, d), (s', d'). Denote this set of source-destination pairs as $\mathcal{D}$.

Lemma 2 implies (via a simple geometric exercise) that all shortest paths connecting (s, d) lie entirely in $T(s, d, 0.5 \times n^{-1/4})$. Thus, no shortest paths connecting different S-D pairs in $\mathcal{D}$ are adjacent. Since a CBR consists of a S-D pair and all shortest paths connecting the source to the destination, the S-D pairs in $\mathcal{D}$ generate mutually non-interfering CBRs. ■

Since paths are closely approximated by line segments in a sufficiently dense random geometric graph, Lemma 1 suggests that a network operating under any PO-MAC can support at most $\Theta(\sqrt{n})$ simultaneous flows. Thus, we conjecture that the user capacity of a network scales as $\Theta(\sqrt{n})$. To verify this conjecture, one would need to show that the maximum set of mutually non-interfering routes selected using any metric – i.e., *non*-shortest paths – cannot scale faster than $O(\sqrt{n})$.

Theorem 3: Consider a set of $n/2$ randomly chosen distinct S-D pairs in dense random geometric graph. With probability tending to 1 as $n \to \infty$, there exist $\Omega(\sqrt{n/\log n})$ S-D pairs which generate mutually non-interfering CBRs. Moreover, these S-D pairs have average distance separation $\Theta(1)$.

Proof: Repeat the constructive part of the proof of Lemma 1, except divide Q into horizontal strips of height $1/\sqrt{n}$ and width 1. With probability tending to 1, we can find $\sqrt{n}/2$ distinct strips that contain line segments. By pruning the number of strips and keeping every $(\sqrt{n}r(n))^{th}$ strip, we are left with a set of $\Theta(\sqrt{n/\log n})$ line segments that are $r(n)$-disjoint. Call this set of line segments $\mathcal{U}$. Consider a single line segment $\ell \in \mathcal{U}$, and denote its endpoints in Cartesian form as: (x_1, y_1) and (x_2, y_2). Then $\|\ell\| \geq |x_1 - x_2| := L_\ell$, where the random variable L_ℓ is independent from the event $\{\ell \in \mathcal{U}\}$ which depends only on y_1, y_2. It is readily shown that $\mathbb{E}[L_\ell] = \Theta(1)$, which proves the theorem. ■

If $X_i \in Q$ is the location of node i in a network, then the *transport capacity* $T(n)$ is defined as the supremum of $\sum_{i \neq j} \lambda_{i,j} \|X_i - X_j\|$, where $\{\lambda_{i,j} : 1 \leq i, j \leq n\}$ is an achievable rate vector [9]. Although not explicitly stated in the literature, $T(n) = \Theta(\sqrt{n/\log n})$ for the Gupta-Kumar protocol in which the transmission radius $r(n) = \Theta(\sqrt{\log n/n})$ is identical for all nodes. To see that $T(n) = O(\sqrt{n/\log n})$, one can repeat the argument in Section 2.5 of [13], and note that the maximum number of simultaneous transmissions is $\Theta(1/r(n)^2)$ instead of $n/2$. Theorem 3 gives the lower bound and shows that the CBR protocol achieves the transport capacity in the sense of Gupta and Kumar.

References

[1] R. Ramanathan, "Challenges: A radically new architecture for next generation mobile ad hoc networks," in *Proc. ACM/IEEE Int'l Conf. Mobile Comp. and Networking*, Cologne, Germany, August 2005.

[2] J. K. Burbank, P. F. Chimento, B. K. Haberman, and W. T. Kasch, "Key challenges of military tactical networking and the elusive promise of MANET technology," *IEEE Comms. Mag.*, pp. 39–45, November 2006.

[3] J. Andrews, N. Jindal, M. Haenggi, R. Berry, S. Jafar, D. Guo, S. Shakkottai, R. Heath, M. Neely, S. Weber, and A. Yener, "Rethinking information theory for mobile ad hoc networks," *IEEE Comms. Mag.*, vol. 46, no. 12, pp. 94–101, December 2008.

[4] A. Blair, T. Brown, K. M. Chugg, and M. J. Johnson, "Tactical mobile mesh network system design," in *Proc. IEEE Military Comm. Conf.*, Orlando, FL, October 2007.

[5] T. R. Halford and K. M. Chugg, "Barrage relay networks," in *Proc. Inform. Theory and Apps. Workshop*, La Jolla, CA, February 2010.

[6] A. Blair, T. Brown, K. M. Chugg, T. R. Halford, and M. Johnson, "Barrage relay networks for cooperative routing in tactical MANETs," in *Proc. IEEE Military Comm. Conf.*, San Diego, CA, November 2008.

[7] T. R. Halford, K. M. Chugg, and A. Polydoros, "Barrage relay networks: System and protocol design," in *Proc. IEEE Int'l Symp. on Personal, Indoor, and Mobile Radio Comms.*, Istanbul, Turkey, September 2010.

[8] T. R. Halford and K. M. Chugg, "The stability of multihop transport with autonomous cooperation," in *Proc. IEEE Military Comm. Conf.*, Baltimore, MD, November 2011.

[9] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.

[10] T. R. Halford and K. M. Chugg, "The reliability of multihop reliability with autonomous cooperation," in *Proc. Inform. Theory and Apps. Workshop*, La Jolla, CA, February 2011.

[11] M. D. Penrose, "The longest edge of the random minimal spanning tree," *Annals of Applied Probability*, vol. 7, no. 2, pp. 340–361, 1997.

[12] P. Gupta and P. R. Kumar, *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W. H. Fleming*. Boston, MA: Birkhauser, 1998, ch. Critical power for asymptotic connectivity in wireless networks, pp. 547–566, eds. W. M. McEneaney, *et al.*.

[13] F. Xue and P. R. Kumar, "Scaling laws for ad hoc wireless networks: An information theoretic approach," *Foundations and Trends in Networking*, vol. 1, no. 2, pp. 145–270, 2006.

[14] N. Alon and J. Spencer, *The Probabilistic Method*. New York: Wiley-Interscience, 2000.

[15] M. Sharir and E. Welzl, "On the number of crossing-free matchings, cycles, and partitions," *SIAM J. Computing*, vol. 36, no. 3, pp. 695–720, 2006.

[16] R. Ellis, J. Martin, and C. Yan, "Random geometric graph diameter in the unit ball," *Algorithmica*, vol. 47, pp. 421–438, 2007.