

Longest edge routing on the spatial Aloha graph

Steven Weber*, Nihar Jindal†, Radha Krishna Ganti°, Martin Haenggi°

* Drexel University, Dept. of ECE, Philadelphia, PA 19104

† University of Minnesota, Dept. of ECE, Minneapolis, MN 55455

° University of Notre Dame, Dept. of EE, Notre Dame, IN 46556

Abstract—The multihop spatial reuse Aloha (MSR-Aloha) protocol was recently introduced by Baccelli *et al.*, where each transmitter selects the receiver among its feasible next hops that maximizes the forward progress of the head of line packet towards its final destination. They identify the optimal medium access probability (MAP) that maximizes the spatial density of progress, defined as the product of the spatial intensity of attempted transmissions times the average per-hop progress of each packet towards its destination. We propose a variant called longest edge routing where each transmitter selects its longest feasible edge, and then identifies a packet in its backlog whose next hop is the associated receiver. The main contribution of this work (and of Baccelli *et al.*) is the use of stochastic geometry to identify the optimal MAP and the corresponding optimal spatial density of progress.

I. INTRODUCTION

A recent paper by Baccelli, Błaszczyszyn, and Mühlethaler [1] introduced the multihop spatial reuse Aloha (MSR-Aloha) protocol in the context of an ad hoc network. The main idea is to combine an Aloha medium access control (MAC) protocol with a routing protocol that moves packets along hops that maximize the progress of each packet towards its final destination. Baccelli *et al.* derive the optimal medium access probability (MAP) that maximizes the spatial density of progress, defined as the number of transmissions per square meter times the average progress towards the destination of each transmitted packet. A key contribution of their paper is the use of stochastic geometry to explicitly incorporate the effect of node locations on network performance.

In this paper we analyze two related but distinct protocols termed random edge routing and longest edge routing. A key assumption in [1] is that each transmitter (Tx) selects the next hop receiver (Rx) as the node that carries the *head of line packet* at the Tx furthest towards its eventual destination, selected over all Rx such that the received signal to interference plus noise ratio (SINR) is sufficiently high to ensure a successful reception. In contrast, we consider a regime where each Tx has a sufficient backlog of packets to ensure it has at least one packet in queue for each potential next hop Rx. Under random edge routing, each Tx selects one of its feasible next hop Rx at random, and selects a packet from its queue appropriate for that Rx. Under longest edge routing, each Tx selects the Rx furthest away. Both protocols are measured under the same spatial density of progress employed in [1].

Our motivation for introducing these two protocols is both practical and mathematical. Practically, the performance of MSR-Aloha, which restricts attention to the head of line packet will by construction be inferior to the performance of longest edge routing, which exploits “buffer diversity” to allow the MAC protocol to move the packets the furthest distance. Mathematically, it is perhaps intuitive to see that analysis of random edge and longest edge routing is simpler than analysis of MSR-Aloha, on account of the fact that progress in the former is measured by the average and maximum edge length, while progress in the latter involves a projection of that length onto the line connecting each packet with its final destination.

There are many practical challenges that can be raised against both MSR-Aloha and our proposed edge routing protocols. Mobility, for example, must be slow enough so that the network topology is sufficiently static to allow a routing protocol to inform all packets of the locations of their final destinations, and by extension their next hop. Our random and longest edge protocols assume a packet backlog covering all possible next hops, but this assumption raises questions about the stability of the queues in the network, and the effect of the protocol on delay. Timing and synchronization are also ignored. While these are valid criticisms, the focus of this paper is on obtaining closed form expressions for the optimal MAP for random edge and longest edge routing, as well as expressions for the spatial intensity of progress under the optimal MAP. The tractability of the model is necessarily reduced if it is extended to incorporate the drawbacks mentioned above.

The rest of this paper is organized as follows. §II introduces the mathematical model. §III presents analytical results on the optimal MAP and corresponding optimal spatial density of progress under random edge and longest edge routing. §IV presents simulation results and shows a good match with the analytical results. A brief conclusion is given in §V. The proofs are placed in the Appendix.

II. MATHEMATICAL MODEL

Consider an infinitely large ad hoc network where the node locations at some snapshot in time form a stationary Poisson point process (PPP) $\Pi = \{x_i\}$ on the plane of intensity λ . During each time slot each node elects to transmit (Tx) with probability p or receive (Rx) with probability $1 - p$; it follows that the set of Tx’s (Π_{Tx}) and Rx’s (Π_{Rx}) are themselves stationary PPPs of intensities λp and $\lambda(1 - p)$ respectively, with $\Pi_{Tx} \cup \Pi_{Rx} = \Pi$. The success of an attempted communication from a Tx to a Rx depends upon the signal to interference plus noise ratio (SINR), measured at the Rx.

This work is supported by the DARPA IT-MANET program, Grant W911NF-07-1-0028 (Jindal, Haenggi, Weber). Corresponding author: Steven Weber, sweber@ece.drexel.edu.

Definition 1: The SINR from each Tx i to each Rx j is

$$\text{SINR}_{ij} = \frac{h_{ij}d_{ij}^{-\alpha}}{\sum_{k \in \Pi_{\text{Tx}} \setminus \{i\}} h_{kj}d_{kj}^{-\alpha} + \eta}, i \in \Pi_{\text{Tx}}, j \in \Pi_{\text{Rx}}, \quad (1)$$

where $\alpha > 2$ is the pathloss exponent, $d_{ij} = d(x_i, x_j)$ is the distance from i to j , $\{h_{kj}\} \sim \text{Exp}(1)$ are the iid Rayleigh fading channel gains, and η is the noise power.

A. MAC: the spatial Aloha graph

We assume a Tx-Rx pair are successful iff the SINR at the receiver exceeds the SINR threshold, β . The spatial Aloha graph was introduced by Ganti and Haenggi in [2].

Definition 2: The spatial Aloha graph is an infinite random geometric directed bipartite graph $G = (\Pi_{\text{Tx}}, \Pi_{\text{Rx}}, E)$, where edges indicate a sufficiently high SINR:

$$(i, j) \in E \Leftrightarrow \text{SINR}_{ij} \geq \beta, i \in \Pi_{\text{Tx}}, j \in \Pi_{\text{Rx}}. \quad (2)$$

A new realization of this graph is created in each time slot, when each node independently decides to Tx or Rx.

Assumption 1: The SINR threshold, β , required for successful reception, is assumed to equal or exceed unity: $\beta \geq 1$. This ensures each Rx has an in-degree of either zero or one.

The in-degree bound follows from the assumption because reception for $\beta > 1$ requires one node have a signal contribution exceeding the sum interference contribution of all other Tx's, and there can necessarily be at most one such node. A realization of this graph is shown in Fig. 1 for $p \in \{0.05, 0.20, 0.50\}$. Qualitative properties of G include:

- **Small p :** There are few Tx's and many Rx's. There is low interference, and hence longer edges. The Tx's each have high out-degree, many Rx's have an in-degree of one.
- **Large p :** More Tx's and fewer Rx's. There is higher interference, and hence only shorter edges are possible. The Tx's each have low out-degree, many Tx in fact have zero out-degree and many Rx have zero in-degree.

The key tradeoff is that although there are many long edges for low p , each Tx can make use of only one of them, hence we wish to have more Tx's (higher p), but the additional Tx's cause more interference, which reduces the number of edges and the average length per edge.

B. Routing: selecting an outgoing edge from each Tx

Coordination is assumed so that in each time slot each Tx knows those Rx's for which the SINR is sufficiently large, and hence knows the set of potential Rx's for that time slot. Given that a Tx may have multiple Rx's, we specify below some possible rules for each Tx to select among the various Rx's.

- **MSR-Aloha** [1]: each Tx selects the Rx maximizing the progress of the head of line packet towards its destination. The spatial density of progress is the intensity of successful Tx times the average packet progress towards its destination.
- **Random / longest edge routing:** each Tx selects a Rx at random (random edge routing), or selects the Rx that is furthest away (longest edge routing), and then selects a packet whose assigned next hop is that Rx. The spatial density of progress is the product of the intensity of successful Tx times the average (or average maximum) edge length.

A ‘‘fair’’ comparison of random/longest edge routing with MSR-Aloha requires selecting a destination for each packet in the buffer, then selecting the packet and Rx pair with the longest progress towards destination, and then using this projected length in evaluating the spatial density of progress. This would spoil the model tractability while gaining little in insight, and so we simply focus on computing the average/maximum edge length. That is, we don't consider the effect of finite buffers, nor do we project the edge lengths onto the line towards each packet's destination. Alternately, note the unprojected edge length is an increasingly accurate measure of progress as the buffer length grows large.

III. OPTIMAL MAP AND SPATIAL DENSITY OF PROGRESS

Assumption 2: Throughout §III we assume that there is no noise, $\eta = 0$, and indicate this by writing SIR instead of SINR.

This assumption is realistic in the interference-limited case. In the present case of Rayleigh fading, a noise term contributes an independent exponential term to the probability of transmission success [1]. The no-noise assumption is unrealistic for very small p , since edges of unbounded length are possible.

A. Preliminary results

We first summarize key results from [1], [2], [3]. First, Baccelli *et al.* [1] established the probability of an edge between a Tx and a Rx separated by distance d .

Proposition 1: (from [1]). The probability that a Tx-Rx pair (i, j) separated by distance d_{ij} has sufficiently high SIR, and hence has an edge in E is

$$\mathbb{P}((i, j) \in E) = \mathbb{P}(\text{SIR}_{ij} \geq \beta) = \exp\{-\pi d_{ij}^2 \lambda p \kappa\}, \quad (3)$$

where

$$\kappa = (\pi\delta) \csc(\pi\delta)\beta^\delta, \quad \delta = 2/\alpha. \quad (4)$$

It is worth noting that the proof of the above result relies critically on the assumed Rayleigh fading; for a general fading distribution (including no fading) one must resort to bounds on the above probability, see *e.g.*, [4]. We next define the neighbors and degrees for each node.

Definition 3: The random set of Rx for each Tx, and the random set (of maximum cardinality one) of Tx for each Rx:

$$\begin{aligned} \mathcal{M}_i^{\text{out}} &= \{j \in \Pi_{\text{Rx}} : (i, j) \in E\}, i \in \Pi_{\text{Tx}} \\ \mathcal{M}_j^{\text{in}} &= \{i \in \Pi_{\text{Tx}} : (i, j) \in E\}, j \in \Pi_{\text{Rx}}. \end{aligned} \quad (5)$$

The random Tx out-degree and random Rx in-degree:

$$M_i^{\text{out}} = |\mathcal{M}_i^{\text{out}}|, i \in \Pi_{\text{Tx}}, \quad M_j^{\text{in}} = |\mathcal{M}_j^{\text{in}}|, j \in \Pi_{\text{Rx}}. \quad (6)$$

The mean out-degree and in-degree

$$m^{\text{out}} = \mathbb{E}[M^{\text{out}}], \quad m^{\text{in}} = \mathbb{E}[M^{\text{in}}] \quad (7)$$

are the expected out (in) degree obtained by selecting a Tx (Rx) uniformly at random over the set Π_{Tx} (Π_{Rx}).

The mean degrees are given by Ganti and Haenggi [2], reproduced below.

Proposition 2: (from [2]). The mean degrees are:

$$m^{\text{out}} = \frac{1-p}{p} \frac{1}{\kappa}, \quad m^{\text{in}} = \frac{1}{\kappa}. \quad (8)$$

It is noteworthy that the mean in-degree is independent of λ and p . The number of edge heads and tails must match; this is seen by weighing the average out (in) degree with the spatial intensity of Tx (Rx): $\lambda p m^{\text{out}} = \lambda(1-p)m^{\text{in}}$. Further, since M_j^{in} is Bernoulli, we in fact know the distribution is

$$\mathbb{P}(M_j^{\text{in}} = 1) = 1 - \mathbb{P}(M_j^{\text{in}} = 0) = m^{\text{in}}, \quad j \in \Pi_{\text{Rx}}. \quad (9)$$

A lower bound (via Jensen's inequality) on the probability of no outgoing edges is given by Ganti and Haenggi [3], reproduced below. Note the bound is independent of λ .

Proposition 3: (from [3]). The probability of no outgoing edges is lower bounded by:

$$\mathbb{P}(M_{\text{out}} = 0) \geq e^{-m^{\text{out}}}. \quad (10)$$

These three propositions are used in the proofs of our main results: Theorem 1 (2) on random (longest) edge routing.

B. Random edge routing (RER)

Define the rv $L = d_{ij}$ as the length of an edge (i, j) selected uniformly at random from the set E in G .

Theorem 1: The average edge length under RER is

$$\mathbb{E}[L] = \frac{1}{2} \sqrt{\frac{m_{\text{in}}}{\lambda p}}. \quad (11)$$

The spatial density of progress is upper bounded by

$$h(\lambda, p) \equiv \lambda p \mathbb{P}(M_{\text{out}} > 0) \mathbb{E}[L] \leq \frac{1}{2} \sqrt{\lambda p m_{\text{in}}} (1 - e^{-m^{\text{out}}}). \quad (12)$$

The bound-optimum MAP is

$$p^* = 2m_{\text{in}} \left(-1 - 2\mathcal{W}_{-1} \left(-\frac{1}{2} e^{-\left(\frac{1}{2} + m_{\text{in}}\right)} \right) \right)^{-1}, \quad (13)$$

where $\mathcal{W}_k(x)$ is the k^{th} branch of the Lambert function, defined as the solution of $\mathcal{W}(x)e^{\mathcal{W}(x)} = x$.

The proof is found in the Appendix. Note that the optimal MAP depends only on κ and not on λ .

C. Longest edge routing (LER)

Define the rv L^{max} as the maximum edge length emanating from a Tx selected uniformly at random from those Tx with one or more Rx in G , i.e., over the set $\{i \in \Pi_{\text{Tx}} : M_i^{\text{out}} > 0\}$.

Theorem 2: The complementary cumulative distribution function (CCDF) of the maximum edge length under LER is approximately

$$\mathbb{P}(L^{\text{max}} > l) \approx \frac{1 - \exp \left\{ -m_{\text{out}} e^{-\frac{\pi l^2 \lambda p}{m_{\text{in}}}} \right\}}{1 - e^{-m^{\text{out}}}}. \quad (14)$$

The average maximum edge length is approximately

$$\mathbb{E}[L^{\text{max}}] \approx \frac{\int_0^\infty \left(1 - \exp \left\{ -m_{\text{out}} e^{-\frac{\pi l^2 \lambda p}{m_{\text{in}}}} \right\} \right) dl}{1 - e^{-m^{\text{out}}}}. \quad (15)$$

The spatial density of progress is approximately

$$h(\lambda, p) \approx \lambda p \int_0^\infty \left(1 - \exp \left\{ -m_{\text{out}} e^{-\frac{\pi l^2 \lambda p}{m_{\text{in}}}} \right\} \right) dl. \quad (16)$$

The proof is found in the Appendix. The approximation in (14) is not necessarily a lower or upper bound since both the numerator and denominator in (26) are upper bounded.

IV. NUMERICAL AND SIMULATION RESULTS

Figures 2, 3, and 4 present numerical and simulation results from Theorems 1 and 2. Simulation results were obtained by taking a Monte-Carlo average over 5 independent realizations of a network arena \mathcal{A} of size 400×400 square meters, with an intensity $\lambda = 0.02$. The average number of nodes was therefore $\mathbb{E}[N] = \lambda|\mathcal{A}| = 3200$. Figures 2 and 3 were obtained using $\alpha = 3$, $\beta = 1$, $\eta = 10^{-6}$, yielding $\kappa \approx 2.4184$ (4).

Fig. 2 shows simulation and numerical results for spatial density of progress, $h(\lambda, p)$, for RER (12) and LER (16), versus the MAP p . The approximate $h(\lambda, p)$ is seen to be quite accurate over all p . The optimal LER achieves 25% higher progress than optimal RER, with 33% fewer attempted Tx.

Fig. 3 shows simulation and numerical results for expected edge length, $\mathbb{E}[L], \mathbb{E}[L^{\text{max}}]$, for RER (11) and LER (15), versus the MAP p . The approximations are quite accurate over all p aside from p near 0. The numerical edge lengths are unbounded as $p \rightarrow 0$ due to the no noise assumption.

Fig. 4 shows numerical results for the optimal medium access probability, p^* , for RER (13) and LER, versus κ (4). The p^* for LER is found by numerically maximizing (16):

$$p^{\text{LER},*} = \arg \max_{p \in [0,1]} \lambda p \int_0^\infty \left(1 - \exp \left\{ -m_{\text{out}} e^{-\frac{\pi l^2 \lambda p}{m_{\text{in}}}} \right\} \right) dl. \quad (17)$$

The inset shows a plot of κ versus the SINR requirement β for pathloss exponents $\alpha = \{2.5, 3, 4, 5\}$. The inset shows κ is increasing in β and decreasing in α , and that $\kappa \geq 1$. The y -axis for the inset, $[1, 20]$, is used as the x -axis for the main figure. For $\alpha = 3$ and $\beta = 1$ ($\kappa \approx 2.4184$), we have $p^{\text{LER},*} \approx 0.14$ and $p^{\text{RER},*} \approx 0.21$, a 50% increase in the optimal MAP.

V. CONCLUSION

Theorems 1 and 2 give approximate (yet very accurate) explicit expressions for the average edge length, the spatial density of progress, and the optimal MAP for RER and LER. The MSR-Aloha protocol proposed in [1] is equivalent to our proposed RER protocol if we ignore projections of link edges onto the final destination line, i.e., if the whole edge length is counted as progress. Our results quantify the improvement of LER over RER both in terms of increased spatial density of progress and in terms of reduced optimal medium access probability. This improvement can be thought of as the Aloha-MAC benefit of exploiting buffer diversity instead of only considering the head of line packet.

REFERENCES

- [1] F. Baccelli, B. Błaszczyszyn, and P. Mühlethaler, "An Aloha protocol for multihop mobile wireless networks," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 421–436, Feb. 2006.
- [2] R. K. Ganti and M. Haenggi, "Dynamic connectivity and packet propagation delay in ALOHA wireless networks," in *Proceedings of the 2007 Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Nov. 2007.
- [3] —, "Single-hop connectivity in interference-limited hybrid wireless networks," in *Proceedings of the 2007 IEEE International Symposium on Information Theory (ISIT)*, Nice, France, June 2007.
- [4] S. Weber, J. G. Andrews, and N. Jindal, "The effect of fading, channel inversion, and threshold scheduling on ad hoc networks," *IEEE Transactions on Information Theory*, vol. 53, no. 11, pp. 4127–4149, Nov. 2007.
- [5] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, 2nd ed. Springer, 2003.




Fig. 1. A realization of the spatial Aloha graph for $p = 0.05$ (top), $p = 0.20$ (middle), and $p = 0.50$ (bottom). Transmitters are denoted by \bullet , receivers by \circ , and an edge indicates an SINR above $\beta = 1$. The square arena has a side length of $200\sqrt{2}$ meters and there are approximately 1600 nodes. The pathloss exponent is $\alpha = 3$ and the noise power is $\eta = 10^{-6}$. Zooming in on the bottom plot reveals many very short edges.

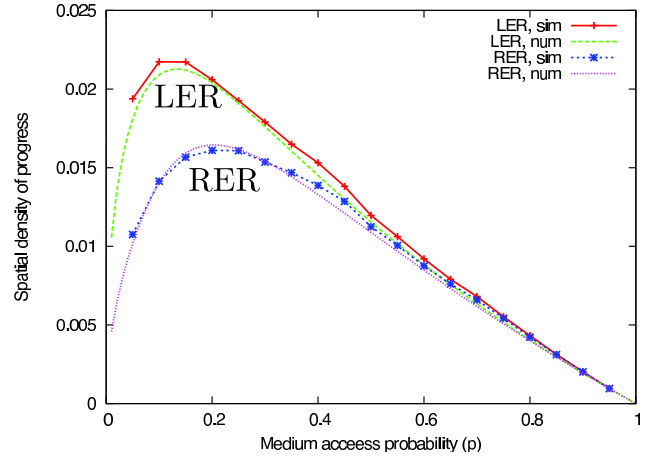


Fig. 2. Simulation and numerical results for spatial density of progress, $h(\lambda, p)$, for random edge routing (RER (12)) and longest edge routing (LER (16)), versus the medium access probability p . Optimal LER achieves 25% higher progress than optimal RER, with 33% fewer attempted transmissions.

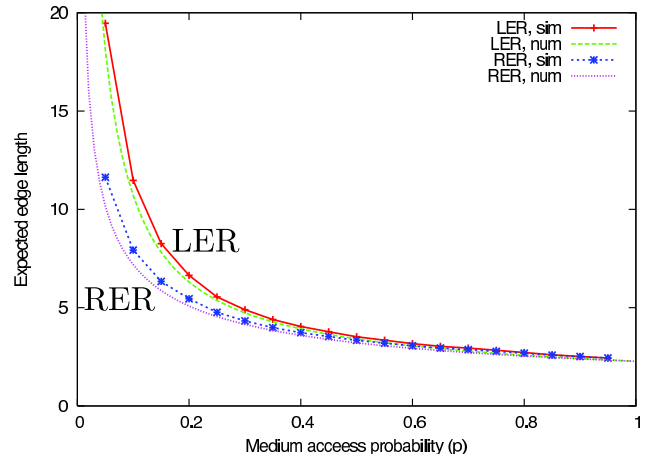


Fig. 3. Simulation and numerical results for expected edge length, $\mathbb{E}[L], \mathbb{E}[L^{\max}]$, for random edge routing (RER (11)) and longest edge routing (LER (15)), versus the medium access probability p . The numerical edge lengths are unbounded as $p \rightarrow 0$ on account of the no noise assumption.

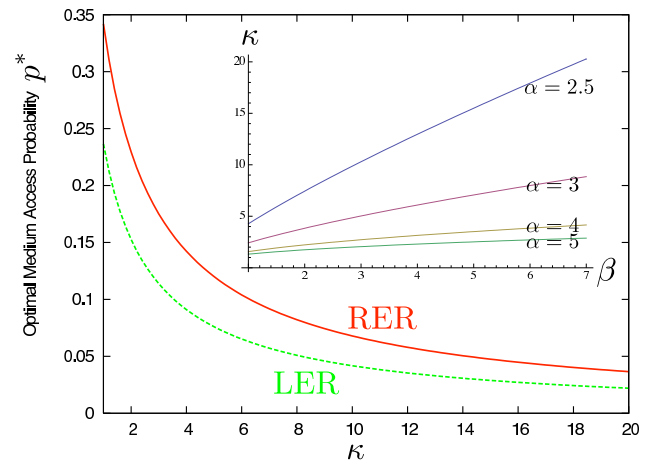


Fig. 4. Numerical results for the optimal medium access probability, p^* , for random edge routing (RER (13)) and longest edge routing (LER), versus κ (4). **Inset:** κ versus the SINR requirement β for pathloss exponents $\alpha = \{2.5, 3, 4, 5\}$.

APPENDIX

PROOF OF THEOREM 1

Proof of (11). Let L_0 be the length of an edge found by selecting a Rx at random. Precisely, L_0 is the length of the edge associated with the Rx, if any is present, or is zero else. Let L be the length of an edge found by selecting an edge at random. We relate $\mathbb{E}[L_0]$ and $\mathbb{E}[L]$ by conditioning: $\mathbb{E}[L_0] =$

$$\mathbb{E}[L_0|M_{\text{in}} = 0]\mathbb{P}(M_{\text{in}} = 0) + \mathbb{E}[L_0|M_{\text{in}} = 1]\mathbb{P}(M_{\text{in}} = 1), \quad (18)$$

so that

$$\mathbb{E}[L] = \mathbb{E}[L_0|M_{\text{in}} = 1] = \frac{1}{m_{\text{in}}}\mathbb{E}[L_0] = \kappa\mathbb{E}[L_0]. \quad (19)$$

Consider a typical Rx located at the origin. Define a *mark* on each Tx $M_i = \frac{\text{SIR}_{i0}}{|x_i|^{-\alpha}}$, let $\{M_i\}$ be the dependent but identically distributed marks for each $i \in \Pi_{\text{Tx}}$, and define F_M as the marginal CDF for a typical mark M . Write L_0 in terms of the marked point process $\Pi'_{\text{Tx}} = \{(x_i, M_i), i \in \Pi_{\text{Tx}}\}$:

$$L_0 = \sum_{i \in \Pi'_{\text{Tx}}} |x_i| \mathbf{1}_E(i, 0) = \sum_{i \in \Pi'_{\text{Tx}}} |x_i| \mathbf{1}_{M_i |x_i|^{-\alpha} > \beta}, \quad (20)$$

where $\mathbf{1}_E(i, 0)$ is the indicator that edge $(i, 0) \in E$. Now apply Campbell's Theorem [5]:

$$\begin{aligned} \mathbb{E}[L_0] &= \lambda p \int_{\mathbb{R}^2} \int_0^\infty |x| \mathbf{1}_{m|x|^{-\alpha} > \beta} dF_M(m) dx \\ &= \lambda p \int_{\mathbb{R}^2} |x| \mathbb{P}(M|x|^{-\alpha} > \beta) dx \\ &= \lambda p \int_{\mathbb{R}^2} |x| \mathbb{P}(\text{SIR}_{0,x} > \beta) dx \\ &= \lambda p \int_{\mathbb{R}^2} |x| e^{-\pi|x|^2 \lambda p \kappa} dx = \frac{1}{2\sqrt{\lambda p \kappa}^{\frac{3}{2}}}. \end{aligned} \quad (21)$$

Proof of (12). The upper bound on the spatial density of progress is obtained by multiplying the intensity of attempted transmissions, λp , times the upper bound on the probability that each transmitter has at least one edge, $\mathbb{P}(M_{\text{out}} > 0) \leq 1 - \exp\{-m_{\text{out}}\}$, times the average length of an edge, $\mathbb{E}[L] = 1/(2\sqrt{\lambda p \kappa})$.

Proof of (13). Maximization of (12) wrt p is equivalent to maximization of

$$h(p) = 2\sqrt{\frac{\kappa}{\lambda}} h(\lambda, p) = \sqrt{p} \left(1 - e^{-\frac{1-p}{\kappa p}}\right). \quad (22)$$

This function has derivative:

$$\sqrt{p} h'(p) = \frac{1}{2} - \left(\frac{1}{2} + \frac{1}{\kappa p}\right) e^{-\frac{1-p}{\kappa p}}. \quad (23)$$

Solving $h'(p) = 0$ for p yields (13).

PROOF OF THEOREM 2

Let $L_0^{\text{max}} \geq 0$ be the maximum edge length of those edges (if any) emanating from a Tx selected at random over all Π_{Tx} . Let this Tx be labeled 0, and wlog, located at the origin, o . Letting $\mathcal{M}_0^{\text{out}}$ be the set of Rx for this Tx, we have:

$$L_0^{\text{max}} = \max_{i \in \mathcal{M}_0^{\text{out}}} |x_i|, \quad (24)$$

with the convention that $L_0^{\text{max}} = 0$ if the set $\mathcal{M}_0^{\text{out}}$ is empty. Let $L^{\text{max}} > 0$ be the maximum edge length of those edges emanating from a Tx selected at random over the set of Tx with one or more edges: $\{i \in \Pi_{\text{Tx}} : M_i^{\text{out}} > 0\}$. We relate the CCDFs of L_0^{max} and L^{max} by conditioning on the random out-degree, $M_0^{\text{out}} \geq 0$ of the Tx selected uniformly over Π_{Tx} :

$$\mathbb{P}(L_0^{\text{max}} > l) = \mathbb{P}(L_0^{\text{max}} > l | M_0^{\text{out}} > 0) \mathbb{P}(M_0^{\text{out}} > 0), \quad (25)$$

since $\mathbb{P}(L_0^{\text{max}} > l | M_0^{\text{out}} = 0) = 0$ for all $l > 0$. Then:

$$\mathbb{P}(L^{\text{max}} > l) = \mathbb{P}(L_0^{\text{max}} > l | M_0^{\text{out}} > 0) = \frac{\mathbb{P}(L_0^{\text{max}} > l)}{\mathbb{P}(M_0^{\text{out}} > 0)}. \quad (26)$$

Define the *marked PPP* $\Pi'_{\text{Rx}} = \{(x_j, \text{SIR}_{0j}), j \in \Pi_{\text{Rx}}\}$, where the marks are the SIRs from Tx 0 to each Rx j . The marks determine whether each pair $(0, j)$ is an edge in E : $(0, j) \in E \Leftrightarrow \text{SIR}_{0j} > \beta$. Note the event equivalence:

$$\begin{aligned} \{L_0^{\text{max}} \leq l\} &= \{x \in B(o, l), \forall x \in \mathcal{M}_0^{\text{out}}\} \\ &= \{\mathbf{1}_E(0, j) \mathbf{1}_{B^c(o, l)}(x_j) = 0, \forall j \in \Pi'_{\text{Rx}}\} \\ &= \left\{ \prod_{j \in \Pi'_{\text{Rx}}} (1 - \mathbf{1}_E(0, j) \mathbf{1}_{B^c(o, l)}(x_j)) = 1 \right\} \end{aligned} \quad (27)$$

for $B(o, l) = \{x \in \mathbb{R}^2 : d_{o,x} \leq l\}$ the ball of radius l centered at the origin, and $B^c(o, l)$ its complement. In words, the event that the maximum edge length for the Tx at the origin is less than l is the same as the event that there are no edges in E from o to receivers outside $B(o, l)$, which is the same as the event that for each receiver j either the pair $(0, j)$ is not in E , or the receiver is in $B(o, l)$. The CDF, $F_0(l) = \mathbb{P}(L_0^{\text{max}} \leq l)$ can be expressed as an expectation:

$$\begin{aligned} F_0(l) &= \mathbb{P} \left(\prod_{j \in \Pi'_{\text{Rx}}} (1 - \mathbf{1}_E(0, j) \mathbf{1}_{B^c(o, l)}(x_j)) = 1 \right) \\ &= \mathbb{E} \left[\prod_{j \in \Pi'_{\text{Rx}}} (1 - \mathbf{1}_E(0, j) \mathbf{1}_{B^c(o, l)}(x_j)) \right]. \end{aligned} \quad (28)$$

This latter expression matches the form required to apply the probability generating function (pgfl) for a (bounded) functional of a PPP ([5]):

$$F_0(l) = \mathbb{E} \left[\exp \left\{ -\lambda(1-p) \int_{\mathbb{R}^2} \mathbf{1}_E(0, x) \mathbf{1}_{B^c(o, l)}(x) dx \right\} \right], \quad (29)$$

where the expectation is wrt the marks $\{\text{SIR}_{0j}, j \in \Pi_{\text{Rx}}\}$. Jensen's inequality yields a lower bound on the CDF:

$$F_0(l) \geq \exp \left\{ -\lambda(1-p) \int_{\mathbb{R}^2} \mathbb{P}(\text{SIR}_{0,x} > \beta) \mathbf{1}_{B^c(o, l)}(x) dx \right\}. \quad (30)$$

Applying (3) and simplifying yields

$$F_0(l) \geq \exp \left\{ -m_{\text{out}} e^{-\frac{\pi l^2 \lambda p}{m_{\text{in}}}} \right\}. \quad (31)$$

Substituting (31) and (10) into (26) yields (14). One finds the approximation for the average maximum edge length (15) by integrating the CCDF (14). Finally, the approximation for the spatial density of progress (16) is obtained by multiplying the same three quantities as in (12); see the Proof of Theorem 1.