Throughput Scaling Laws for Wireless Ad Hoc Networks with Relay Selection

Invited Paper

Marios Kountouris and Jeffrey G. Andrews
Wireless Networking and Communications Group
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX 78712-0240
Email: {mkountouris, jandrews}@ece.utexas.edu

***Abstract*—We consider transmission of packets in two-hop wireless ad hoc networks in which relay nodes are deployed between the source-destination pairs. Based on results from extreme value theory and product tails, we derive throughput scaling laws when opportunistic relay selection is performed. Assuming partial channel state information at each transmitter (CSIT) and decode-and-forward, half-duplex relays, we investigate how the per-hop throughput depends on the channel gain asymptotic distribution and the relay deployment. In dense networks with λ_t nodes per m^2 and fixed relay distances, we provide specific scaling laws for Rayleigh, lognormal, and Weibull fading, showing that the throughput is upper bounded by $\Theta(\sqrt{\lambda_t})$. Interestingly, with variable relay distances and location-aware relay selection, we analytically show that regularly varying channel distributions result in enhanced multi-relay diversity gain, achieving linear throughput scaling $\Theta(\lambda_t)$.**

I. Introduction

Determining the capacity of uncoordinated wireless ad hoc networks is one of the most general and challenging problems in network information theory. The current mainstream approach consists of deriving sum rate bounds and asymptotic capacity scaling laws, which describe important aspects of the capacity region. In their seminal work [1], Gupta and Kumar studied the communication limits in an ad hoc wireless network where n nodes are uniformly distributed in a plane and grouped into source-destination (S-D) pairs at random. It was shown that employing multihop nearest-neighbor routing, a sum throughput of $\Theta(\sqrt{n/\log n})$ can be achieved. Using percolation theory, it has been shown that in a network with random S-D pairs, a partially-connected multihop routing strategy can achieve capacity of order $\Theta(\sqrt{n})$ [2]. Actually, in spite of special cases that suggest more optimistic scalings [3], [4], it is generally agreed that the scaling law in ad hoc networks is upper bounded by $\Theta(\sqrt{n})$ [2], [5]–[7].

Two network models are commonly used in order to derive asymptotic in n scaling laws: *dense networks*, in which a growing number of nodes are distributed over a fixed area, and *extended networks* where a fixed density of nodes is placed over a growing area. Following this line of research, which often provides useful high-level insights how different network configurations may affect the scaling law, we study the effect of opportunistic relay selection on the capacity of random spatial networks which are well-modeled by a Poisson field of interference. Note that a dense network can be constructed by placing nodes according to a Poisson point process (PPP) of intensity λ over a square of unit area. Similarly, we can construct an extended network by placing nodes according to a PPP of unit intensity on the plane and focus our attention to the square $[0, \sqrt{\lambda}] \times [0, \sqrt{\lambda}]$, so that the average number of nodes is $\lambda = n$.

The main focus of this paper is on relay-assisted communications with feedback, and specifically on studying how to exploit multi-relay diversity gains in two-hop wireless ad hoc networks. Feedback-based relay selection has been studied in [8] and [9], where a CSI-based opportunistic scheduling protocol is proposed. However, the effect of pathloss attenuation is neglected in both works. In [10], a simple distributed relay selection method is proposed exploiting channel reciprocity.

In this work, consider decentralized relay selection protocols, where every node makes its next-hop decision based only on local channel state information (CSI), namely its channel gain and distance from the transmitting nodes. We investigate the throughput scaling of two-hop dense networks in which opportunistic relay selection is performed under two different relay deployment strategies: deterministic and random. Specifically, assuming that each receiver/relay has perfect CSI, whereas sources rely on partial CSI conveyed via low-rate feedback, we provide a general upper bound on the performance of such networks, showing that the per-hop throughput is upper bounded by $O(\sqrt{\lambda})$ (order optimal). Furthermore, since the exact throughput scaling depends on the specific class of the channel distribution, we show that the capacity scales as $\Theta(\log \lambda)$ for Rayleigh fading, $\Theta(e^{\sqrt{2 \log \lambda}\sigma})$ for lognormal fading, and $\Theta(\log^{1/k} \lambda)$ for Weibull fading, assuming that all candidate relays are placed at equal distances from the sources. In contrast, when the relay nodes are randomly distributed around each transmitter, a linear throughput scaling $\Theta(\lambda)$ can be achieved. This significant multi-relay diversity gain is due to the heavy tail behavior of the resulting total received power. More generally, if the total received power is distributed according to a Pareto-type distribution (power law), the received signal power is asymptotically of

This research has been supported by the DARPA IT-MANET program.

the order of magnitude of the resulting aggregate interference when nearest neighbor routing is performed.

II. Network Model

We consider a network in which the transmitting nodes (sources) are independently and randomly distributed according to a stationary Poisson point process Φ_t with intensity λ_t on the plane. A set of relays Φ_r of intensity λ_r is also deployed between the source and destination pairs. We assume that every source has a destination (not a part of Φ_t or Φ_r) at a distance D_m it wishes to send packets to. The relays are neither sources nor destinations, and two different deployment strategies are studied: (1) *deterministic*, where the relays are placed at fixed distances νD and τD from each source and destination respectively, with $D_m > D$ and $\nu, \tau \in (0,1)$, and (2) *random*, where the relays are located at random in a finite disk around the sources. For exposition convenience, $\nu = \tau = 1$ is assumed in the remainder, without reducing the validity of the derived scaling laws. Since the relays do not necessarily form a PPP, some care must be taken to ensure that this model is valid. Refining this two-hop spatial reuse model and relay selection analysis for inhomogeneous Poisson point processes, where the relays belong to a thinned Φ_t, is left for future work.

A. Channel Model

The existence of a communication link from source node i to node j depends on the ability of the receiver j to decode the transmitted signal, which is in turn determined by the signal-to-interference ratio (SIR) given by

$$\mathrm{SIR}_{ij} = \frac{\rho S_{ij} \ell(D_{ij})}{\sum_{k \in \mathcal{S}_t, k \neq i} \rho S_{kj} \ell(D_{kj})} \quad (1)$$

where $\mathcal{S}_t$ is the transmitting set, ρ is the fixed transmission power, and S_{ij} is the channel fading coefficient between nodes i and j. The fading S results from independent and identically distributed (i.i.d.) fading distributions, and, unless otherwise stated, is assumed to have finite mean and variance. The pathloss attenuation function $\ell(\cdot)$ is assumed to depend only on the distance D_{ij} between node i and j, and is chosen to be $\ell(D_{ij}) = D_{ij}^{-\alpha}$, where $\alpha > 2$ is the pathloss exponent. Noise can also be included, but it has a negligible effect for interference-limited random networks and asymptotic capacity analysis, which is our case of interest.

B. Key Performance Metrics

A primary performance metric of interest in an uncoordinated, random ad hoc network is the success probability $\mathcal{P}_{\mathrm{succ}}$, which is defined as the probability that the SIR at a typical S-D pair is above a pre-defined target SIR β. A message is successfully decoded if $\mathcal{P}_{\mathrm{succ}} = \mathbb{P}\{\mathrm{SIR} \geq \beta\}$, where the SIR statistics are a function of the interferers density $|\mathcal{S}_t| \leq \lambda_t$ and $\mathcal{P}_{\mathrm{succ}}$ is clearly a decreasing function of λ_t.

Assuming transmission at the Shannon target rate $R = \log_2(1+\beta)$ bps/Hz, the per-hop throughput (area spectral efficiency) is defined as

$$\mathcal{C}_\epsilon \triangleq \mathcal{N}_t \mathcal{P}_{\mathrm{succ}} R \quad \text{bps/Hz} \quad (2)$$

where $\mathcal{N}_t$ is the number of concurrent transmission attempts.

Notation: We denote $\mu_X = \mathbb{E}(X)$ and $\sigma_X^2 = \mathrm{Var}(\mathrm{X})$ the expected value and the variance of a random variable X, respectively. For any functions $f(n)$ and $g(n)$, $g(n) = O(f(n))$ is equivalent to $\lim_{n\to\infty} |g(n)/f(n)| < \infty$, $g(n) = \Theta(f(n))$ implies that $\lim_{n\to\infty} |g(n)/f(n)| = c$, where $0 < c < \infty$, and $g(n) \sim f(n)$ is equivalent to $\lim_{n\to\infty} g(n)/f(n) = 1$.

III. Two-Hop Opportunistic Relaying

In this section, we derive an upper bound on the throughput of two-hop ad hoc networks when opportunistic relaying is performed. For that, general classes of channel fading and interference distributions are considered, assuming only the existence of finite mean and variance.

A. Opportunistic Relay Scheduling

We adopt an opportunistic relay selection scheme (for the first hop) where, similarly to multiuser diversity, only source-relay pairs with favorable conditions are activated. We assume that each receiver/relay has perfect channel state information (CSI), whereas transmitters rely on (local) partial CSI conveyed via low-rate feedback channel. Specifically, a relay k selects its preferred transmitter node i^*, as the one with the highest signal strength Z_{ik}, *i.e.*

$$i^* = \arg\max_{i \in \Phi_t} Z_{ik} = \arg\max_{i \in \Phi_t} S_{ik} D_{ik}^{-\alpha}. \quad (3)$$

Each relay reports the index i^* to the corresponding transmitter, resulting in a total feedback load of λ_r integer values. The set of selected relays $\mathcal{S}_r$ has cardinality $|\mathcal{S}_r| \leq \lambda_r$, where the inequality is strict if a transmitter is served by multiple relays.

B. Upper Bound on the Success Probability

First, we derive a general upper bound on the probability that $|\mathcal{S}_t|$ communication pairs can communicate at rate $R = \log_2(1+\beta)$ bps/Hz. Let $S_{\max} = \max_{1 \leq i \leq |\mathcal{S}_t|} S_{ik}$ denote the channel fading between a relay k and its preferred source i^*, which corresponds to the maximum value from a sample of $|\mathcal{S}_t|$ i.i.d. random variables (r.v.) with common parent distribution $F_S(\cdot)$. Assuming that S has finite mean variance and that the resulting interference I_Φ has finite moments, we can state that

***Lemma** 1: The probability of successful communication between a relay and a source selected by means of multi-relay diversity is given by*

$$\mathcal{P}_{\mathrm{succ}} \leq \frac{\mu_{S_{\max}}}{\beta D^\alpha} \mathbb{E}\left\{\frac{1}{I_\Phi}\right\} \leq \frac{\mu_S + \frac{|\mathcal{S}_t|-1}{\sqrt{2|\mathcal{S}_t|-1}}\sigma_S}{\beta D^\alpha} \mathbb{E}\left\{\frac{1}{I_\Phi}\right\}.$$

Proof: See Appendix A. ■

The interference term I_Φ is the summation of $(|\mathcal{S}_t|-1)$ channel gains conditioned that the received signal is $S_{\max}$, thus its exact characterization is difficult due to its dependence with S (and the inhomogeneity and thinning of the original PPP). Note that the aggregate interference is a shot-noise process whose distribution can be efficiently approximated by a gamma or inverse Gaussian distribution [11] and can be shown to asymptotically converge to a stable distribution [3]. We also point out that the aforementioned lemma is not restricted to the PPP assumption, and can be applied to a variety of interference models. Consider a dense random network where $\mathbb{E}\left\{I_\Phi^{-1}\right\} = O(\lambda_r^{-c})$ with $0 \leq c \leq 1$ and $|\mathcal{S}_r| = \lambda_r = O(\lambda_t^b)$. It can be easily shown that for $b = 1/(2c)$, the success probability is $\mathbb{P}\{\mathrm{SIR} \geq \beta\} = O(1)$. Thus, the throughput scaling will depend on the growth of the expected number of successful concurrent transmissions $\mathcal{N}_t$.

For $c = 1$, we see that channel-aware opportunistic relay selection results in a success probability of order $O(1)$, in contrast to single-hop communications in which the success probability is upper bounded by $O(\lambda_t^{-1})$ since the received signal growth is upper bounded by $O(\sqrt{\lambda_t})$ and the interference scales as $O(\lambda_t)$. As the number of source-relay pairs is upper-bounded by $|\mathcal{S}_r|$, setting $\lambda_r = \frac{\lambda_t - 1}{\sqrt{2\lambda_t - 1}} = O(\sqrt{\lambda_t})$ and exploiting multi-relay diversity gain, the throughput scaling of the first hop is upper bounded by $O(\sqrt{\lambda_t})$.

In the second hop, a scheme that guarantees $\mathbb{P}\{\mathrm{SIR} \geq \beta\} > 0$ for some target β and all $|\mathcal{S}_r|$ relay-destination communication links is required, as a means to preserve the capacity scaling of the first hop. In contrast to single-hop communications and since delay analysis is out of the scope of this work, each relay (of unbounded buffer) can decode the message of any transmitting source and forward it to the intended receiver when the channel conditions are favorable.

***Proposition** 1: The throughput of the second-hop under opportunistic selection, in which each destination node m schedules any relay $k \in \Phi_r$ for which the received* $\mathrm{SIR} \geq \beta$, *is given by*

$$\mathcal{C} = |\mathcal{S}_r| \left(1 - \mathbb{P}\{\mathrm{SIR}_{km} \leq \beta\}\right) R = |\mathcal{S}_r| \left(1 - F_{\mathrm{SIR}}^{\lambda_d}(\beta)\right) R$$

and scales as $\mathcal{C} = \Theta(\sqrt{\lambda_d})$ for $\lambda_d \to \infty$, where $\lambda_d = \lambda_t$ is the intensity of destination nodes.

It can be easily shown that such scheme results in $\mathcal{P}_{\mathrm{succ}} = O(1)$. For instance, if the relay-destination pairs form an homogeneous PPP with intensity λ_r, we have that [12]

$$F_{\mathrm{SIR}}(\beta) = 1 - \exp\left\{-\lambda_r \pi D^2 \beta^{\frac{2}{\alpha}} \frac{2\pi}{\alpha} \csc\left(\frac{2\pi}{\alpha}\right)\right\}.$$

Thus, for $\lambda_t \to \infty$ (thus $\lambda_r \to \infty$) we have $\left(1 - F_{\mathrm{SIR}}^{\lambda_t}(\beta)\right) \to 1$, which results in capacity of order $O(\sqrt{\lambda_t})$. Note also that the interference of the first hop stochastically dominates the second-hop interference, since the latter is the sum of $|\mathcal{S}_r|$ terms excluding the $S_{\max}$. Finally, since the two-hop throughput is determined by the minimum achievable rate of each hop, we conclude that the above two-hop random networks with relay selection can achieve a throughput of $\Theta(\sqrt{\lambda_t})$.

IV. Scaling Laws with Multi-Relay Diversity

In the previous section, we provided order-wise throughput performance results based on a general upper bound on the success probability. Nevertheless, the exact capacity scaling depends on the specific type of distribution to which signal strength belongs to. In this section, we provide capacity scaling laws by characterizing the resulting limiting distribution of S due to opportunistic scheduling. For that, we take on a simple way to exploit partial CSI and adopt a channel threshold scheduling policy, under which the relay elects a source provided the channel gain exceeds a specified threshold [13]. Although we claim no optimality, eliminating the fading coefficients below the threshold can significantly reduce outage for many fading distributions of interest, and therefore can significantly increase the intensity of transmissions. Let $\mathcal{A}_\gamma$ be the event that the channel strength between a selected transmitter-relay pair is above a threshold γ, i.e. $S_{i^*k} \geq \gamma$. Note that threshold-based scheduling changes the distribution of S_{i^*k} (for all i) from the unconditional distribution S to the conditional distribution of S given $S \geq \gamma$, but leaves the distribution of S_{jk} for $j \neq i^*$ unchanged.

***Proposition** 2: Under threshold-based scheduling and opportunistic relay selection, the success probability is lower bounded by*

$$\mathbb{P}\{\mathrm{SIR} \geq \beta\} \geq (1 - F_{S_{\max}}(\gamma)) \frac{(\frac{\gamma D^{-\alpha}}{\beta} - \mu_{I_\Phi})^2}{\sigma_{I_\Phi}^2 + (\frac{\gamma D^{-\alpha}}{\beta} - \mu_{I_\Phi})^2} \quad (4)$$

when the mean and the variance of the interference are finite, and

$$\mathbb{P}\{\mathrm{SIR} \geq \beta\} \geq (1 - F_{S_{\max}}(\gamma)) \left(1 - \frac{\mu_{I_\Phi}}{\gamma} \beta D^\alpha\right) \quad (5)$$

when the interference has finite mean.

Proof: See Appendix B. ■

The first term of the product determines the probability that the scheduled transmitter's channel gain exceeds a threshold γ and the second term (RHS) gives the probability that the SIR between source and relay is above β.

Based on the above lower bound, we aim at deriving the throughput scaling laws in dense networks ($\lambda_t \to \infty$), as well as the optimal relay density λ_r for specific types of fading distributions. Note that the lower bound is valid even for finite λ_t, λ_r. We distinguish between two different relay deployments: the first model assumes that all relays are a fixed distance away from their transmitting nodes. The second model considers random relay deployment and a variable distance between relay and source.

A. Deterministic Relay Deployment

The capacity scaling depends on the asymptotic scaling of success probability and the number of scheduled sources. Due to opportunistic relay selection, success probability scaling depends in turn on the domain of attraction of the channel fading distribution and the asymptotic scaling of the interference of CDF. For dense ad hoc networks with large λ_t and λ_r, the following lemma can be shown

***Lemma** 2: The capacity of the first hop with opportunistic relay selection scales as*

For Rayleigh fading:	$\mathcal{C} = \Theta(\log \lambda_t)$
For lognormal fading:	$\mathcal{C} = \Theta(e^{\sigma_S \sqrt{2 \log \lambda_t}})$
For Weibull fading with shape k:	$\mathcal{C} = \Theta(\log^{1/k} \lambda_t)$

Proof: See Appendix C. ■

In Fig. 1 we plot the average throughput of a two-hop wireless ad hoc network as a function of node intensity λ_t for $\beta = 1$, $\alpha = 4$ and shape $k = 1.4$ for Weibull fading. As expected, the throughput under lognormal fading is considerably higher than that obtained with exponential fading, while Weibull fading result in the lowest performance.

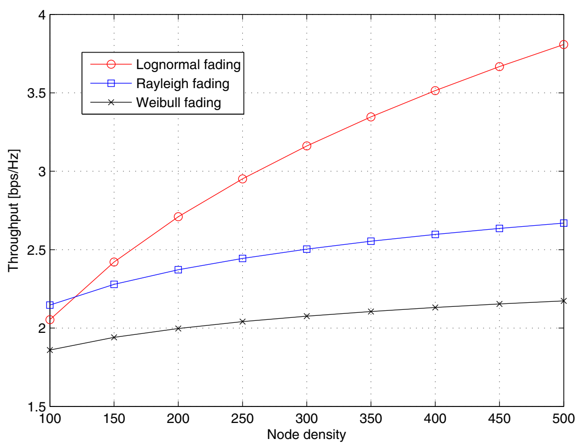

Fig. 1. Throughput of two-hop opportunistic relaying vs. node intensity λ_t for $\alpha = 4$ and $\beta = 1$.

B. Random Relay Deployment

In the second model, we consider a more practical approach and assume that the transmission relay distances are variable. We take on a simplified scenario in which the transmitter is located at the center of a circular area of radius $D_{\max}$ and its relay nodes are uniformly distributed over the area. Therefore, the inter-node distance D_{ik} is a random variable with non uniform distribution $f_D(d)$ with

$$f_D(d) = \frac{2d}{D_{\max}^2 - D_{\min}^2}, \quad \text{for } d \in [D_{\min}, D_{\max}]. \tag{6}$$

Assuming that the relays are dropped randomly in each area, the random process D_{ik} can be considered i.i.d. across relays and transmitters. Note also that edge effects of the disk can be neglected in the regime we are interested in. Relays falling outside the circular area can be dropped from the asymptotic analysis without affecting the scaling law.

For exposition convenience, we assume $D_{\max} - D_{\min} = 1$m. It can be easily shown that the distribution (CDF) of pathloss attenuation $X_{ik} = D_{ik}^{-\alpha}$ is given by

$$F_X(x) = 1 - x^{-\frac{2}{\alpha}}, \quad x \geq 1. \tag{7}$$

Remarkably, the distribution in (7) differs from typical fast fading distributions (e.g. Rayleigh, Nakagami, Rice) as it exhibits *heavy tail* behavior, which in turn plays a fundamental role in shaping the limiting distribution of its maximum value.

***Proposition** 3: The distribution $F_X(x)$ is a* regularly varying *function (at ∞) with exponent $-2/\alpha$,* i.e.

$$\frac{1 - F_X(x)}{1 - F_X(ux)} \to u^{2/\alpha} \quad \textit{when} \quad u \to \infty. \tag{8}$$

To characterize the tail behavior of the product $Z = SX$ (assuming that channel fading distribution S has finite moments), we exploit an interesting aspect of regularly varying distributions that states that they are stable with respect to multiplication with other independent r.v. with finite moments. Applying Breiman's theorem [14] we have

***Lemma** 3: Let X be a r.v. with distribution given by (7) and S be an independent r.v. with $\mathbb{E}(S^{\frac{2}{\alpha}}) < \infty$. Then the tail of $Z = SX$ is governed by*

$$1 - F_Z(x) \sim \mathbb{E}(S^{\frac{2}{\alpha}}) x^{-\frac{2}{\alpha}} \quad \textit{when } x \to \infty. \tag{9}$$

Therefore, the total channel power Z is a regularly varying distribution with exponent $-2/\alpha$, as the heavy tail behavior of the distribution of random relay distances tends to dominate over the channel fading distribution. We can now state our main result:

***Theorem** 1: In two-hop random wireless networks with variable relay distances, location-aware opportunistic relay selection results in two-hop throughput that scales linearly with the node density, i.e. $\mathcal{C} = \Theta(\lambda_t)$.*

Proof: See Appendix D. ∎

The above theorem implies that when the total channel power Z follows a heavy tail distribution (Pareto type), the signal strength at the nearest neighbor is asymptotically of the order of magnitude of the aggregate interference, thus the expected density of successful concurrent transmissions may be $\mathcal{N} = \lambda_t \mathbb{P}\{\text{SIR} > \beta\} = \Theta(\lambda_t)$. Thus, deploying $\lambda_r \geq \lambda_t$ relays and selecting the nearest node results in linear throughput scaling for $\lambda_t \to \infty$.

Similar linear scaling has been reported in [3], due to the fact that node mobility results in a Pareto distributed received power. Nevertheless, only large-scale pathloss attenuation is considered in [3], whereas our proposed framework generalizes previous linear throughput results as it may also include multipath and shadowing effects.

V. Conclusions

We have examined the asymptotic throughput scaling of large random wireless ad hoc networks with feedback-assisted opportunistic relaying. Our results show that two-hop relay communications using deterministically deployed relays is order optimal, while when the relays are distributed randomly around the transmitters, a linear throughput scaling can be achieved in dense networks. Location-aware scheduling can fully exploit the multi-relay diversity benefits, providing significant performance gains in two-hop ad hoc communications.

Appendix

A. Proof of Lemma 1

For non-negative random interference I_Φ with finite moments we have

$$\mathbb{P}\{\text{SIR} \geq \beta\} = \mathbb{P}\left\{\frac{S_{\max}}{I_\Phi} > \beta D^\alpha\right\}$$
$$\stackrel{(a)}{=} \int_0^\infty \mathbb{P}\{S_{\max} > x\beta D^\alpha\} f_{I_\Phi} dx \stackrel{(b)}{\leq} \int_0^\infty \frac{\mu_{S_{\max}}}{x\beta D^\alpha} f_{I_\Phi} dx$$
$$= \frac{\mu_{S_{\max}}}{\beta D^\alpha} \mathbb{E}\left\{\frac{1}{I_\Phi}\right\} \stackrel{(c)}{\leq} \frac{\mu_S + \frac{\lambda_t - 1}{\sqrt{2\lambda_t - 1}} \sigma_S}{\beta D^\alpha} \mathbb{E}\left\{\frac{1}{I_\Phi}\right\}$$

where (a) is reached by conditioning on $I_\Phi = x$, the Markov inequality is used for (b), and an upper bound on the expected value of the highest order statistic is used for (c) [15].

B. Proof of Proposition 2

For a predefined threshold γ, the total probability theorem results in the following lower bound on the success probability:

$$\begin{aligned}\mathbb{P}\{\mathrm{SIR} \geq \beta\} &= \mathbb{P}\{\mathcal{A}_\gamma\} \cdot \mathbb{P}\{\mathrm{SIR} \geq \beta | \mathcal{A}_\gamma\} \\ &+ \mathbb{P}\{\mathcal{A}_\gamma^c\} \cdot \mathbb{P}\{\mathrm{SIR} \geq \beta | \mathcal{A}_\gamma^c\} \\ &\geq \mathbb{P}\{\mathcal{A}_\gamma\} \cdot \mathbb{P}\{\mathrm{SIR} \geq \beta | \mathcal{A}_\gamma\} \\ &\geq (1 - F_{S_{\max}}(\gamma)) \cdot \mathbb{P}\{I_\Phi < \gamma D^{-\alpha}/\beta\}\end{aligned}$$

To derive (4), the Chebyshev-Cantelli inequality is employed, which gives that $\mathbb{P}\{I_\Phi \geq \gamma D^{-\alpha}/\beta\} \leq \frac{\sigma^2_{I_\Phi}}{\sigma^2_{I_\Phi} + (\gamma D^{-\alpha}/\beta - \mu_{I_\Phi})^2}$, while for (5) we apply the Markov inequality that states $\mathbb{P}\{I_\Phi \geq \gamma D^{-\alpha}/\beta\} \leq \frac{\mu_{I_\Phi}}{\gamma}\beta D^\alpha$.

C. Proof of Lemma 2

The throughput of the first hop is lower bounded by $\mathcal{C} \geq \mathcal{N}_r \mathbb{P}\{\mathrm{SIR} \geq \beta\} R \geq \mathcal{N}_r (1 - F_{S_{\max}}(\gamma)) \left(1 - \frac{\mu_{I_\Phi}}{\gamma}\beta D^\alpha\right) R$, where $\mathcal{N}_r$ is a lower bound on the number of selected sources. Each source has a probability of $1/\lambda_t$ to be selected by a relay, thus λ_r distinct source-relay pairs are formed with probability $\mathcal{P} = \frac{\lambda_t(\lambda_t - 1)\ldots(\lambda_t - \lambda_r + 1)}{\lambda_t^{\lambda_r}}$.

The proof is composed by three parts: We first derive the domain of attraction $\mathcal{D}_i(x)$ of the channel power for different channel distributions. For that, we need to characterize the normalizing sequences a_{λ_t}, b_{λ_t} such that

$$1 - F_S(b_{\lambda_t}) = \frac{1}{\lambda_t} \quad \text{and} \quad a_{\lambda_t} = h(b_{\lambda_t}) \tag{10}$$

so that $\mathbb{P}\left\{\frac{S_{\max} - b_{\lambda_t}}{a_{\lambda_t}} \leq x\right\} \to \mathcal{D}_i(x)$, with $i = 1, 2, 3$ denoting the Gumbel, Fréchet, and Weibull domain of attraction, respectively. The function $h(x) = \frac{1 - F_S(x)}{f_S(x)}$ is the so-called hazard rate.

- For Rayleigh fading, since $h^{'}(x) = 0$ for $x \to \infty$, $S_{\max}$ converges to $\mathcal{D}_1$ with sequences $a_{\lambda_t} = 1$ and $b_{\lambda_t} = \log \lambda_t$. Selecting a threshold $t = \log \lambda_t - \log\log \lambda_t$, we have $F_{S_{\max}}(\gamma) = (1 - e^{-\gamma})^{\lambda_t} \to 0$.
- For lognormal fading, we have that $h(x) = \frac{\sqrt{2}\sigma_S x}{2}\left(\sum_{i=0}^{\infty}(-1)^i \frac{(2i)!}{i!(2x)^{2i}}\right)$ (using the asymptotic expansion of complementary error function), which implies that lognormal distribution converges asymptotically to Gumbel domain of attraction (since $\lim_{x\to\infty} h^{'}(x) = 0$). For normalizing sequences $a_{\lambda_t} = \frac{b_{\lambda_t}\sigma_S}{\sqrt{2\log\lambda_t}}$ and $b_{\lambda_t} = e^{\sqrt{2\log\lambda_t}\sigma_S + \mu_S}$ and threshold $\gamma = e^{\sigma_S\sqrt{2\log\lambda_t} + \mu_S - e^{\sigma_S\sqrt{2\log\lambda_t} + \mu_S}}$, we have that $1 - F_{S_{\max}}(\gamma) = \Theta(1)$.
- For Weibull fading with scale $k_1 > 0$ and shape $k_2 > 0$, we have that $h(x) = \frac{(1-k_2)k_1^{k_2}}{k_2}x^{-k_2}$, thus with $a_{\lambda_t} = k_1/k_2$ and $b_{\lambda_t} = k_1(\log\lambda_t)^{1/k_2}$ the channel gain distribution converges to $\mathcal{D}_1$. For $\gamma = (\log\lambda_t - \log\log\lambda_t)^{1/k_2}$, we have $F_{S_{\max}}(\gamma) = (1 - e^{-(\gamma/k_1)^{k_2}})^{\lambda_t} \to 0$, thus $1 - F_{S_{\max}}(\gamma) = \Theta(1)$.

Secondly, for all three cases, the expected density of source-relay pairs approaches $\mathcal{N}_r \to \lambda_r$ for $\lambda_r = \Theta(\log\lambda_t)$, $\lambda_r = \Theta(e^{\sigma_S\sqrt{2\log\lambda_t}})$, and $\lambda_r = \Theta((\log\lambda_t)^{1/k_2})$, respectively. Finally, for the intereference term, we have that $1 - \frac{\mu_{I_\Phi}}{\gamma}\beta D^\alpha = \Theta(1)$ for $\lambda_r = \log\lambda_t$ (Rayleigh), $\lambda_r = \Theta(e^{\sqrt{2\log\lambda_t}\sigma_S + \mu_S})$ (lognormal), and $\lambda_r = k_2(\log\lambda_t)^{1/k_2}$ (Weibull), respectively.

D. Proof of Theorem 1

Invoking Smirnov's theorem, we can show that the limiting distribution of $Z_{\max}$ is of Fréchet type with

$$\mathbb{P}\left\{\left(\frac{\kappa}{\log^2\sqrt{\lambda_t}}\right)^{\frac{\alpha}{2}} \leq Z_{\max} \leq \kappa^{\frac{\alpha}{2}}\right\} \geq 1 - O\left(\frac{1}{\log\lambda_t}\right) \tag{11}$$

with $\kappa = \lambda_t \mathbb{E}(S^{\frac{2}{\alpha}})\log\sqrt{\lambda_t}$. Using arguments from [3] and the Continuous Mapping Theorem, we can show that $Z_{\max}$ and I_Φ are asymptotically independent, and that the interference asymptotically converges to a stable distribution with characteristic exponent $2/\alpha$. Since asymptotically $\mathbb{E}\{Z_{\max}/I_\Phi\} \to 1 - 2/\alpha$ [16], we have $\mathbb{P}\{\mathrm{SIR} \geq \beta\} > 0$ for $\beta = 1 - 2/\alpha$. This implies that for $\lambda_r = \lambda_t$ relays, the average number of successful concurrent transmissions with rate $R = \log_2(1+\beta)$ bps/Hz is of the order of λ_t, thus $\mathcal{C} = \Theta(\lambda_t)$.

REFERENCES

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. on Info. Theory*, vol. 46, no. 2, pp. 388-404, Mar. 2000.

[2] M. Franceschetti, M. Migliore, and P. Minero, "The capacity of wireless networks: Information-theoretic and physical limits," *submitted to IEEE Trans. on Info. Theory*, available at http://circuit.ucsd.edu/∼massimo/.

[3] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad-hoc wireless networks," *IEEE/ACM Trans. on Networking*, vol. 10, no. 4, pp. 477-486, Aug. 2002.

[4] A. Ozgur, O. Leveque, and D. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," *IEEE Trans. on Info. Theory*, vol. 53, no. 10, pp. 3549-3572, Oct. 2007.

[5] O. Leveque and I. E. Telatar, "Information-theoretic upper bounds on the capacity of large extended ad hoc wireless networks," *IEEE Trans. on Info. Theory*, vol. 51, no. 3, pp. 858-865, Mar. 2005.

[6] 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.

[7] R. Gowaikar, B. Hochwald, and B. Hassibi, "Communication over a wireless network with random connections," *IEEE Trans. on Info. Theory*, vol. 52, no. 7, pp. 2857-2871, Jul. 2006.

[8] C. K. Lo, R. W. Heath Jr., and S. Vishwanath, "Opportunistic relay selection with limited feedback," *Proc. IEEE VTC-Spring'07*, Dublin, Ireland, pp. 135-139, Apr. 2007.

[9] S. Cui, A. M. Haimovich, O. Somekh, and H. V. Poor, "Opportunistic relaying in wireless networks," *submitted to IEEE Trans. on Info. Theory*, available at http://arxiv.org/abs/0712.1169, 2007.

[10] A. Bletsas, A. Lippnian, and D. P. Reed, "A simple distributed method for relay selection in cooperative diversity wireless networks, based on reciprocity and channel measurements," *Proc. IEEE VTC-Spring'05*, Stockholm, Sweden, pp. 1484-1488, May-Jun. 2005.

[11] R. K. Ganti and M. Haenggi, "Interference in ad hoc with general motion-invariant node distributions," *Proc. IEEE Intl. Symposium of Information Theory*, Toronto, Canada, Jul. 2008.

[12] F. Baccelli, B. Blaszczyszyn, and P. Muhlethaler, "An ALOHA protocol for multihop mobile wireless networks," *IEEE Trans. on Info. Theory*, vol. 52, no. 2, pp. 421-436, Feb. 2006.

[13] S. Weber, J. G. Andrews, and N. Jindal, "The effect of fading, channel inversion, and threshold scheduling on ad hoc networks," *IEEE Trans. on Info. Theory*, vol. 53, no. 11, pp. 4127-4149, Nov. 2007.

[14] L. Breiman, "On some limit theorems similar to the arc-sin law," in *Theory of Probability and its applications, 10.*, pp. 323-331, 1965.

[15] H. A. David and H. N. Nagaraja, *Order Statistics,* third ed., New York: John Wiley and Sons, 2003.

[16] N. H. Bingham and J. L. Teugels, "Conditions implying domains of attraction," *Proc. of the 6th Conf. on Prob. Theory, Editura Academiei Republicii Socialiste*, Bucuresti, Romania, pp. 23-34, 1981.