

On the Catalyzing Effect of Randomness on the Per-Flow Throughput in Wireless Networks

Florin Ciucu
University of Warwick

Jens Schmitt
University of Kaiserslautern

Abstract—This paper investigates the throughput capacity of a flow crossing a multi-hop wireless network, whose geometry is characterized by general randomness laws including Uniform, Poisson, Heavy-Tailed distributions for both the nodes' densities and the number of hops. The key contribution is to demonstrate how the per-flow throughput depends on the distribution of 1) the number of nodes N_j inside hops' interference sets, 2) the number of hops K , and 3) the degree of spatial correlations. The randomness in both N_j 's and K is advantageous, i.e., it can yield larger scalings (as large as $\Theta(n)$) than in non-random settings. An interesting consequence is that the per-flow capacity can exhibit the opposite behavior to the network capacity, which was shown to suffer from a logarithmic decrease in the presence of randomness. In turn, spatial correlations along the end-to-end path are detrimental by a logarithmic term.

I. INTRODUCTION

A groundbreaking work at the intersection between communication networks and information theory is a set of network capacity results obtained by Gupta and Kumar [18]. Under some simplifications at the network layers (e.g., no multi-user coding schemes, or ideal assumptions on power control, routing, and scheduling), those results establish asymptotic scaling laws on the maximal data rates which can be reliably sustained in multi-hop wireless networks. They have further inspired a tremendous research interest and provided fundamental insight into the development of the long desirable functional network information theory (see Andrews *et al.* [4]).

A major assumption in the network model from [18], which was widely adopted thereafter, is that the geometry follows a *uniform* distribution. This assumption implies that the nodes' interference sets follow the binomial (or its Poisson approximation) distribution. A small set of works considered geometries with specific heterogeneous geometries, which were shown to play a fundamental role in the capacity scaling laws (see the Related Work Section). The goal of this paper is to go beyond specific random geometries, by analyzing (throughput) capacity results in networks with *general* distributions¹. To make the analysis manageable, the paper assumes that network nodes implement the Aloha MAC protocol; the simplicity of the protocol permits the derivation of per-flow capacity results in *closed-form* and *explicit* up to the optimization of a single parameter.

We point out that the paper particularly focuses on *per-flow capacity* results, rather than the network capacity results which

are mostly treated in the literature. The advantage of per-flow capacity results is that they determine the network capacity by a summation argument; in contrast, a division argument to compute the per-flow capacity from the network capacity only holds for uniform geometries. Moreover, unlike the majority of existing *asymptotic* capacity results—whose practicality is often questioned for small to medium sized networks (Akyildiz and Wang [2], p. 180)—this paper provides *non-asymptotic* results, e.g., the per-flow capacity in (finite) T time units in a network with (finite) N nodes.

Intuitively, different randomness laws in the geometry yield different (per-flow) capacity results. This paper goes beyond this simple intuition and makes three key contributions:

- 1) Analyzing the beneficial role of randomness in the network's geometry on the per-flow capacity.
- 2) Quantifying the magnitude of the benefit, referred to as *randomness gain*, in terms of scaling laws.
- 3) Providing the “best” distributions maximizing the randomness gain.

To properly summarize our observations, let us briefly describe the network model; see Section III for the complete description. A flow crosses some end-to-end (e2e) path with K hops. The interference set of a hop j consists of N_j (interfering) nodes, which is referred to as the hop density; all N_j 's and K can have general distributions. The degree of spatial correlations, denoted by γ , defines the number of consecutive hops with correlated densities; for instance, in the practical scenario when all N_1, N_2, \dots, N_K are statistically correlated, then $\gamma := K$.

The closed-form expressions of the derived capacity results allow a sensitivity study of the various parameters in the network model, e.g., N_j 's, K , and γ . From this study we collected the following observations:

- O1. By scaling N (here a shorthand for N_j 's), the per-flow capacity depends on N 's distribution under a sample-path neighbor-aware probabilities assumption. Concretely, nodes should set transmission probabilities, explicitly or implicitly, according to the nodes' densities (transmission probabilities should be roughly proportional to the number of neighbors). With this assumption, different distributions of N with identical averages can yield different scaling laws. In turn, when nodes use a fixed optimized transmission probability on all sample paths, the per-flow capacity is sensitive to the distribution of N , but only through its mean $E[N]$.

¹In this paper, the notions of throughput and (throughput) capacity, used interchangeably, explicitly refer to the maximal flows' throughput achieved in some network, with no coding schemes being considered. The meaning of *capacity* employed herein differs thus from the one in information theory.

- O2. By scaling time, the per-flow capacity is invariant to K . In other words, if the system runs over a sufficiently long time scale, then the per-flow capacity, with the interpretation of a *rate*, stabilizes and does not depend on the actual number of hops. More interestingly, by simultaneously scaling both time and K , we find that randomness in K has also a fundamental impact on the per-flow capacity (i.e., it changes the order of growth).
- O3. The size of the randomness gain can be as large as $\mathcal{O}(n)$ in N and much smaller in K , which indicates that temporal correlations due to N are much more sensitive to randomness than spatial correlations due to K , as an effect of spatial reuse². Based on specific distributions of N and K , we find the surprising fact that the randomness gain in K is the logarithm of the randomness gain in N . To clarify its precise meaning, the randomness gain is defined as the relative difference of the per-flow capacities in two scenarios: one in which a network parameter (say N) is random, and another in which the same parameter N is set to its (non-random) expectation $E[N]$.
- O4. By simultaneously scaling time, K , and γ , the per-flow capacity decays logarithmically in γ . The simultaneous scaling is needed as in 2); otherwise, the impact of spatial correlation vanishes. The observed logarithmic detrimental factor of spatial correlations, is analogous to an existing result characteristic to *wired* networks: in a tandem of nodes in which exponentially sized packets arrive as a Poisson process, the e2e delay scales as $\Theta(n \log n)$ and not as $\Theta(n)$ [7]. The $\Theta(n)$ holds under the so-called Kleinrock's Independence Assumption that packets independently regenerate their sizes at each hop. Without this assumption, the extra logarithmic factor stems from the spatial correlations due to the "very large" packets inducing long delays to packets behind them, and at *each* node. Similarly, in our settings, the logarithmic term arises when spatial correlations span across the entire e2e path.

These insights complement some fundamental existing ones. One is that randomness can have a detrimental role in the network capacity (see Gupta and Kumar [18] or Franceschetti *et al.* [15]). In contrast, our results show that, in networks with non-uniform geometries, randomness has a beneficial role in the per-flow capacity; for a discussion clarifying the apparently contradicting detrimental and beneficial roles of randomness see the end of Section V-A. Another important known fact is that TDMA and CSMA with properly tuned parameters achieve the same capacity (Chau *et al.* [10]). Such a MAC *insensitivity* result, however, depends on the assumed uniform geometry. In turn, our results indicate that in a single-hop scenario with non-uniform random geometry, the per-flow capacity is in fact *sensitive* to the MAC since CSMA implicitly satisfies the sample-path assumption, whereas explicit overhead would be needed for Aloha or TDMA.

²By *temporal correlations* we mainly refer to dependent events which can occur simultaneously (i.e., transmissions within the same interference set). By *spatial correlations* we mainly refer to dependent events which can occur at different times (i.e., transmissions at two relays whose interference sets share some nodes).

The rest of the paper is organized as follows. First we discuss related work. In Section III, we introduce the network model, and, in Section IV, we introduce the main analytical tools enabling the capacity analysis. In Section V, we first present the main result of the paper, i.e., non-asymptotic bounds on the capacity of a fixed source-destination pair, and then we investigate the capacity's sensitivity to the randomness factors in geometry. Brief conclusions are presented in Section VI.

II. RELATED WORK

Gupta and Kumar [18] analyzed the asymptotic capacity of homogeneous random networks with uniformly distributed nodes, and showed the notorious $\Theta(1/\sqrt{n \log n})$ scaling law on the per-flow capacity under a specific communication channel model. This law was improved to $\Theta(1/\sqrt{n})$ for another channel model by Franceschetti *et al.* [15]. Under a mobility model and a two-hop relay model, the per-flow scaling laws were further improved to $\Theta(1)$, i.e., the best achievable one, but at the expense of conceivably long delays (Grossglauser and Tse [17]). For a more comprehensive review of related scaling laws see Xue and Kumar [36].

Asymptotic capacities for heterogeneous networks, e.g., not necessarily with uniformly distributed nodes, were derived in special cases. Toumpis [32] considered a logically clustered network in which n sources communicate with n^d cluster heads (yet all are uniformly placed), and showed that network capacity degrades in the presence of bottlenecks when $0 < d < 0.5$. Pervalov *et al.* [28] considered physically clustered networks, with uniformly placed nodes and clusters of nodes, and showed that network capacity fundamentally depends on the size of the clusters. For some clustered networks, Kulkarni and Viswanath [24] showed that network capacity preserves the scaling law from [18]. For some other specific clustered networks, however, Alfano *et al.* [3] and Martina *et al.* [26] recently showed that the per-flow capacity is fundamentally influenced by the geometry. Similar results have also been reported from simulations by Hoydis *et al.* [20]. Our paper differs from these works in that it provides per-flow and (non)-asymptotic capacity results for a broad range of random geometries.

As far as non-asymptotic capacity results are concerned, many exist in the single-hop case (e.g., Kleinrock and Tobagi [23] or Bianchi [5]); the latter is derived for 802.11 DCF networks, by assuming that all nodes independently see the system in steady state. Much fewer results exist in multi-hop networks, mostly under simplifying technical assumptions and approximations to deal with the intrinsically hard problem of spatio-temporal correlations. Capacity results were computed in 802.11 DCF networks, modelled as contention graphs, under the assumption that collision probabilities are mainly due to hidden node interference (Gao *et al.* [16]). E2e delays in both TDMA and Aloha line-networks were investigated using a decomposition approach, which relies on the approximation that the departure processes at the relay nodes have independent inter-departure times (Xie and Haenggi [35]). The delay analysis of wireless channels under Markovian assumptions was studied by Zheng *et al.* [37], and the delay analysis of multi-hop fading channels by Al-Zubaidy *et al.* [38].

Closer to our work, non-asymptotic per-flow capacity bounds were derived in networks with *non-random* N_j 's and K , and *no* spatial correlations (i.e., $\gamma = 1$) (see Ciucu *et al.* [12], [11], [13]). This paper extends these results by accounting for general randomness in N_j 's, K , and spatial correlations (i.e., $\gamma > 1$).

Our results on the advantageous effect of randomness relate to a ‘‘folk theorem’’ from queueing theory which states that, when the mean inter-arrival (service) time is fixed, the constant inter-arrival (service) time distribution *minimizes* queueing metrics such as average waiting time. Such results were proven for renewal processes (Rogozin [29]) and also for more general arrival processes with exponential service times (Hajek [19] and Humblet [21]). Moreover, our results on the bimodal nature of distributions maximizing the randomness gain agree with parallel results from queueing theory. For instance, bimodal distributions maximize queue lengths in GI/M/1 queues (Whitt [33]), in G/M/1 queues with bulk arrivals (Lee and Tsitsiklis [25]), or in queues with bulk arrivals and finite buffers (Bušić *et al.* [8]).

III. NETWORK MODEL

In this section, we describe the network model and the type of capacity results investigated in this paper.

We consider a general network model accounting for three randomness sources, thus significantly generalizing related models. Concretely, we consider the multi-hop random network geometry from Fig. 1. Node 1 (the source) transmits to node $K+1$ (the destination) using nodes $2, 3, \dots, K$ as relays, where K is a random variable denoting the number of hops. The number of nodes inside the interference set (IS) of node j , and *excluding* node j , is denoted by the random variable N_j , for $j = 2, 3, \dots, K+1$; N_j 's are also referred to as nodes' (hops') densities. The ISes allow to model arbitrary interference models and do not rely on geometrical assumptions like disc-based transmission or interference ranges.

One requirement is the existence of an e2e path between the source and the destination. This assumption is motivated by the very goal of the paper, i.e., the derivation of per-flow capacities which requires the flow to be well-defined in terms of an e2e path. If such e2e paths were subject to discontinuities, the derived capacity results would still hold for the transient regimes during which an e2e path exists.

The model further needs knowledge of the distributions of N_j 's and K , which characterize the first two randomness sources. These distributions can be quite general; in fact, the capacity formulae from the main result (see Theorem 1) allow plugging-in any specific distribution law in order to quantify the underlying impact.

Concretely, all N_j 's are finite and identically distributed with density

$$\pi_n = \mathbb{P}(N = n), \quad n = 2, \dots, n_{\max}. \quad (1)$$

N generically stands for N_j 's. Note that $\pi_1 = 0$, i.e., the nodes on the e2e path are not isolated, and n_{\max} denotes the maximum number of nodes inside an IS. The assumption of identically distributed N_j 's is a mild one and mitigates notational complexity.

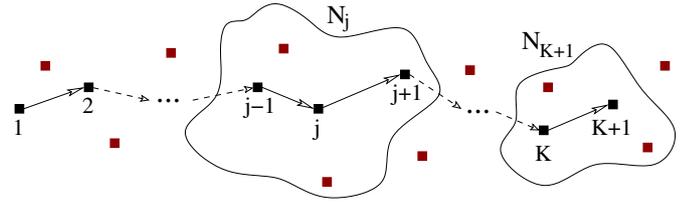


Fig. 1. A multi-hop wireless network with a random number of K hops. The interference set (IS) of node j contains a random number N_j of nodes. The performance metric of interest is the (non-)asymptotic e2e capacity of node 1 transmitting to node $K+1$ using the nodes $2, 3, \dots, K$ as relays.

For the second randomness source, we assume that the number of hops K is finite, statistically independent of all the N_j 's, and has the density

$$\tilde{\pi}_k = \mathbb{P}(K = k), \quad k = 1, 2, \dots, k_{\max}. \quad (2)$$

The maximal number of hops is k_{\max} . The independence assumption simplifies the proof of the main result, i.e., Theorem 1; the proof could also account for the conditional distributions $\mathbb{P}(N_j = n | K = k)$ at the expense, however, of increasing notational complexity. Nevertheless, the independence assumption between K and N_j 's is conceivably strong in small networks. Indeed, if nodes set large interfering ranges then the N_j 's are also large and K is small; in turn, if the interfering ranges are small then the N_j 's are also small and K is large. Such correlation effects lessen in large network regimes, whereby our later asymptotic analysis applies.

The third source of randomness in the network model concerns the degree of spatial correlations, i.e., the degree of statistical dependence amongst N_j 's. Dealing with all possible combinations and types of such dependencies is clearly an overwhelming task. In our analysis we assume that for each $j = 1, 2, \dots, k_{\max}$, the r.v. N_j is statistically independent of all N_i 's with $i \in \{j + \gamma, j + \gamma + 1, \dots, k_{\max}\}$. By the commutativity of the independence relationship between two r.v.'s, N_j is also statistically independent of N_i 's with $i \in \{1, 2, \dots, j - \gamma\}$.

This particular *dependency parameter* γ characterizes the maximal number of consecutive ISes for which dependencies (correlations) may exist between the first and the rest. For instance, if $\gamma = 1$ then all N_j 's are statistically independent; at the other extreme, in a static scenario with k hops, $\gamma = k$ corresponds to dependencies between any pair of N_1, N_2, \dots, N_k (this latter scenario is conceivably the *most practical* one).

Besides the description of the three randomness sources, the network model requires the specification of the MAC protocol. Concretely, given a slotted time model, the network model requires that all nodes transmit with some probability p , independently of each other, and from slot to slot. This requirement is immediately satisfied by the slotted-Aloha protocol (Abramson [1]), and implicitly satisfied by 802.11 DCF under an independence assumption from Bianchi [5], in the single-hop case. The assumption from [5], which is commonly adopted in the literature, states that all nodes see the channel in steady state, while the transmission probability p can be computed from the parameters of the protocol (see Eqs. (7) and (9) in [5]). An additional approximation for 802.11 DCF is the consideration of an average slot size (given

also in [5], in Eq. (13)). In a multi-hop case the steady-state assumption becomes however less justifiable, due to the hidden node problem, and further assumptions are needed (Gao *et al.* [16]). Since we are seeking rigorous capacity results, in order to perform a rigorous analysis of the underlying roles of the three randomness sources, we mainly adopt slotted-Aloha and make side remarks on 802.11 DCF and also on TDMA.

Denoting the transmission from node i to j by $[i \rightarrow j]$, a transmission $[j \rightarrow j+1]$ is successful if node j is the only transmitting node inside the IS of node $j+1$, in a time slot. As far as data sources are concerned, we assume that the nodes $2, 3, \dots, K$ have infinite buffers and only relay the data from node 1 to node $K+1$. Also, all the other nodes are saturated, i.e., they always attempt to transmit according to the MAC protocol. This saturation assumption implies that the computed capacity of the path $[1 \rightarrow K+1]$ is conservative from the data-link layer perspective.

The concrete type of capacity investigated in this paper is the per-flow (non-)asymptotic throughput capacity of the transmission $[1 \rightarrow K+1]$. Denote by $A(t)$ the arrival process at node 1 (i.e., counting the data units to be transmitted to node $K+1$). Also, denote by $D(t)$ the corresponding departure process at node $K+1$ (i.e., counting the data units arrived from node 1 up to time t). For some fixed violation probability ε , a probabilistic upper bound on the (e2e throughput) capacity rate is a value λ_t^U such that

$$\mathbb{P}(D(t) \geq \lambda_t^U t) \leq \varepsilon. \quad (3)$$

In turn, a probabilistic lower bound on the capacity rate is a value λ_t^L such that

$$\mathbb{P}(D(t) \leq \lambda_t^L t) \leq \varepsilon. \quad (4)$$

We point out that the derived capacity rates λ_t^U and λ_t^L are obtained in transient (non-asymptotic in the time scale) regimes, whereas the asymptotic results are immediately obtained by letting $t \rightarrow \infty$. The other non-asymptotic regime is in the nodes' densities N and number of hops K ; again, related asymptotic results follow directly by taking limits.

IV. ANALYTICAL TOOLS

In this section, we introduce the modelling tools for the per-flow capacity problem. The main engine behind the derivations in the paper is the framework of the stochastic network calculus, in particular following Fidler [14]. The advantage of the calculus approach is that it considerably simplifies the complexity of modelling a whole network by (logically) reducing it to a single-hop only. It is thus sufficient to model any single-hop scenario from Fig. 1 with source j and destination $j+1$; the multi-hop model will directly follow from the convolution theorem in the stochastic network calculus.

Let us model an arbitrary single-hop where the source is node j . For each node l inside the IS of node $j+1$ we associate a random process $X_l(t)$, where t represents time. For every l and t , $X_l(t)$ is a Bernoulli random variable (r.v.) taking value 1 with probability p ; with abuse of notation, $l=1$ refers to the source j . These r.v.'s are mutually *i.i.d.* in both time and space, and conditioned on the realizations of the number of

nodes/hops (for a parallel analytical framework, dealing with non-necessarily independent increments of $X_l(t)$, e.g., when modelling CSMA/CA besides Aloha, we refer to Ciucu *et al.* [13]).

Next we introduce the key process for computing the per-flow capacity. This is referred to as the (virtual) interfering process of the transmission $[j \rightarrow j+1]$, and is defined through its increments $V(t-1, t) := V(t) - V(t-1)$ for $t \geq 1$ as

$$V(t-1, t) = 1 - X_1(t) \prod_{l=2}^{N_{j+1}} (1 - X_l(t)). \quad (5)$$

The initial value is $V(0) = 0$. We make the important remark that $V(t)$ does not depend on whether the source j is saturated or bursty, since it is defined independently of the arrival process $A(t)$ at the source. Moreover, as we mentioned earlier, our analysis handles the situation of idle periods characteristic to relay nodes $j \geq 2$, and which are due to internal burstiness: there is nothing to transmit at some slot, and yet the MAC protocol may successfully select the relay node to access the channel. Due to such situations we emphasize the attribute *virtual* for the process $V(t)$.

Next, we obtain the moment generating function (MGF) and the Laplace transform of $V(t)$, needed to derive the upper and lower bounds, respectively, from Eqs. (3) and (4). For some parameter $\theta > 0$, these transforms are defined as

$$M_t(\theta) = E[e^{\theta V(t)}] \quad \text{and} \quad L_t(\theta) = E[e^{-\theta V(t)}].$$

The MGF follows from the backwards equations using conditioning, and also using the independent increments property of $V(t)$, i.e.,

$$\begin{aligned} M_{t+1}(\theta) &= M_t(\theta) E[e^{\theta V(1)}] \\ &= M_t(\theta) \sum_i E[e^{\theta V(1)} | N(0) = i] \mathbb{P}(N(0) = i) \\ &= M_t(\theta) \sum_i (e^\theta q_i + 1 - q_i) \pi_i. \end{aligned}$$

Therefore, by evaluating the sum,

$$M_t(\theta) = b_\theta^t, \quad (6)$$

where

$$b_\theta = 1 + q(e^\theta - 1), \quad q = \sum_{l \geq 2} q_l \pi_l, \quad q_l = 1 - p(1-p)^{l-1},$$

and where the π_l 's are from Eq. (1). The Laplace transform follows by a sign change, i.e.,

$$L_t(\theta) = b_{-\theta}^t.$$

The critical role of the virtual process is to link the arrival process $A(t)$ at the source j with the corresponding departure process $D(t)$ at the destination $j+1$. This relationship is expressed in terms of a stochastic service process, which is a key tool in the stochastic network calculus (Chang [9], Jiang and Liu [22]), and which is instrumental herein for the derivation of e2e capacity results. The next lemma from Ciucu [11] formally establishes this relationship.

Lemma 1: (SINGLE-HOP (EXACT) SERVICE REPRESENTATION) Consider the interfering process $V(t)$ from Eq. (5). Then the bivariate random process

$$S(s, t) = t - s - V(s, t) \quad (7)$$

is an *exact* stochastic service process for node A , i.e.,

$$D(t) = A * S(t) \text{ a.s.}, \quad (8)$$

for all arrival processes $A(t)$. Here, the symbol ‘ $*$ ’ denotes the $(\min, +)$ convolution operator defined for all $t \geq 0$ as $A * S(t) := \inf_{0 \leq s \leq t} \{A(s) + S(s, t)\}$.

The process $S(s, t)$ quantifies the service received over the link $[j \rightarrow j+1]$. A key observation is that Eq. (8) holds for *all* arrival processes $A(t)$. This invariance is instrumental for carrying out the incoming multi-hop analysis, by circumventing the intrinsically difficult (queueing) problem that the arrival processes at the relay nodes $j \geq 2$ are hard to characterize. The process $D(t)$ from Eq. (8) can be also viewed as the output from a variable capacity node (Boudec and Thiran [6]), given the definition of the interfering process $V(t)$ from Eq. (5). Moreover, $S(s, s+1)$ can be viewed as the instantaneous per-flow *effective capacity*, as proposed to model the instantaneous channel capacity by Wu and Negi [34]. In turn, the process $V(s, t)$ can be viewed as an *impairment process* as defined by Jiang and Liu [22], p. 72. Having available an exact service process for single-hop transmissions, we can next derive both upper and lower bounds on e2e capacity.

V. END-TO-END PER-FLOW CAPACITY

In this section, we first derive the main result of this paper, i.e., a closed-form expression for the per-flow capacity along the e2e path from Fig. 1. Then we investigate its sensitivity to the three randomness sources in the network model: the distribution of the number of neighbors N_j 's and the number of hops K , and the dependency parameter γ .

The procedure for getting the upper and lower capacity bounds follows the methodology of the stochastic network calculus (see Boudec and Thiran [6], Chang [9], and Jiang and Liu [22]). First, service processes $S_j(s, t)$ for each single-hop transmission $[j \rightarrow j+1]$, $j = 1, 2, \dots, K$, are constructed as in Eq. (7). These processes are then convolved in the underlying $(\min, +)$ -algebra yielding the (network) service process

$$S(s, t) := S_1 * S_2 * \dots * S_K(s, t), \quad (9)$$

which characterizes the available service (or the per-flow effective capacity in the terminology from Wu and Negi [34]) along the e2e path $[1 \rightarrow K+1]$. Eq. (9) is the convolution theorem from network calculus which reduces the multi-hop analysis to a single-hop analysis. The advantage of the theorem is that it circumvents the difficult problem of modelling input/output processes at intermediate relay nodes, as mentioned earlier.

Theorem 1: (NON-ASYMPTOTIC CAPACITY BOUNDS) Consider the multi-hop network model from Section III with dependency parameter γ . Let $q_l = 1 - p(1-p)^{l-1}$, $q = \sum_{l \geq 2} q_l \pi_l$, and $b_\theta = 1 + q(e^\theta - 1)$ for any $\theta > 0$ (see Eq. (6)). Then, for some violation probability ε , a probabilistic lower bound on the e2e capacity is for all $t \geq k_{max}$

$$\lambda_t^L = \sup_{\theta > 0} \left\{ 1 - \frac{1}{\gamma} \frac{\log b_\theta}{\theta} + \frac{\log \varepsilon}{t\theta} - \frac{c_K}{t\theta} \right\}, \quad (10)$$

where $c_K = \sum_{k=1}^{k_{max}} \tilde{\pi}_k \log \binom{t+k-1}{k-1}$. The upper bound is

$$\lambda_t^U = \inf_{\theta > 0} \left\{ 1 + \frac{1}{\gamma} \frac{\log b_{-\gamma\theta}}{\theta} - \frac{\log \varepsilon}{t\theta} \right\}. \quad (11)$$

We remark that the asymptotic lower and upper bounds coincide (after using Stirling's approximation for the factorial in the binomial term, with $\theta = \Theta(t^{-\zeta})$, $0 < \zeta < 1$), i.e.,

$$\lambda := \lim_{t \rightarrow \infty} \lambda_t^L = \lim_{t \rightarrow \infty} \lambda_t^U = 1 - q. \quad (12)$$

We denote the asymptotic capacity by λ . Theorem 1 generalizes existing non-asymptotic lower-bound results from Ciucu *et al.* [12], [11], [13] which hold for *non-random* N_j 's and K , and $\gamma = 1$. Moreover, Theorem 1 also provides the corresponding upper bounds. We also point out that the results are explicit up to optimizing after $\theta > 0$.

PROOF. Let \mathbb{P}_k denote the underlying probability measure conditioned on $K = k$. Let $t \geq 0$ and the service processes $S_j(s, t)$ for each transmission $[j \rightarrow j+1]$ for $j = 1, 2, \dots, k$, as in Eq. (7). Applying the e2e service process from Eq. (9), we can write

$$\begin{aligned} \mathbb{P}_k(D(t) \leq \lambda_t^L t) &\leq \mathbb{P}_k(A * S(t) \leq \lambda_t^L t) \\ &= \mathbb{P}_k(S_1 * S_2 * \dots * S_k(t) \leq \lambda_t^L t) \end{aligned} \quad (13)$$

because of the saturation condition $A(1) = \infty$. Letting $u_0 = 0$ and $u_k = t$ we can continue above as follows

$$\begin{aligned} &\mathbb{P}_k \left(\inf_{0 \leq u_1 \leq \dots \leq u_{k-1} \leq t} \sum_{j=1}^k S_j(u_{j-1}, u_j) \leq \lambda_t^L t \right) \\ &= \mathbb{P}_k \left(\inf_{0 \leq u_1 \leq \dots \leq u_{k-1} \leq t} \sum_{j=1}^k (u_j - u_{j-1} - V_j(u_{j-1}, u_j)) \leq \lambda_t^L t \right) \\ &= \mathbb{P}_k \left(\sup_{0 \leq u_1 \leq \dots \leq u_{k-1} \leq t} \sum_{j=1}^k (V_j(u_{j-1}, u_j) - (u_j - u_{j-1})) \geq -\lambda_t^L t \right), \end{aligned} \quad (14)$$

where $V_j(u_{j-1}, u_j) = u_j - u_{j-1} - S_j(u_{j-1}, u_j)$. Next, by applying the Union and Chernoff bounds, we can bound the last term in Eq. (14) by

$$\sum_{0 \leq u_1 \leq \dots \leq u_{k-1} \leq t} E \left[\prod_{j=1}^k e^{\theta V_j(u_{j-1}, u_j)} \right] e^{-\theta(1-\lambda_t^L)t}.$$

At this point we rearrange the terms in the product in the expectation as

$$\prod_{j=1}^k e^{\theta V_j(u_{j-1}, u_j)} = \prod_{l=1}^{\gamma} \prod_{i \geq 0} e^{\theta V_{l+i\gamma}(u_{l+i\gamma-1}, u_{l+i\gamma})}. \quad (15)$$

The terms in the second product in the right-hand side term are statistically independent. This is true according to the definition of the dependency parameter γ , and also because the underlying Bernoulli r.v.'s in $V_{l+i\gamma}$'s have independent

increments; note the non-overlapping intervals of $V_{l+i\gamma}$'s. With this observation we can bound the last expectation by

$$\left(\prod_{l=1}^{\gamma} E \left[\prod_{i \geq 0} e^{\gamma \theta V_{l+i\gamma}(u_{l+i\gamma-1}, u_{l+i\gamma})} \right] \right)^{\frac{1}{\gamma}} = (b_{\gamma\theta}^t)^{\frac{1}{\gamma}},$$

by using Hölder's inequality, where $b_{\gamma\theta}$ has the expression from the theorem with θ replaced by $\gamma\theta$. Recall also the MGF of $V(t)$ from Eq. (6).

Collecting terms we obtain

$$\begin{aligned} \mathbb{P}_k(D(t) \leq \lambda_t^L t) &\leq \sum_{0 \leq u_1 \leq \dots \leq u_{k-1} \leq t} (b_{\gamma\theta}^t)^{\frac{1}{\gamma}} e^{-\theta(1-\lambda_t^L)t} \\ &= \binom{t+k-1}{k-1} (b_{\gamma\theta}^t)^{\frac{1}{\gamma}} e^{-\theta(1-\lambda_t^L)t}, \end{aligned}$$

where the binomial term is the number of combinations with repetition.

For the upper bound, we recall from Theorem 1 that the service processes $S_j(s, t)$ are exact, and, therefore, the e2e service process from Eq. (9) is exact as well; this property is critical for proving the upper bound. Using again that $A(1) = \infty$ we can write

$$\begin{aligned} \mathbb{P}_k(D(t) \geq \lambda_t^U t) &= \mathbb{P}_k(A * S(t) \geq \lambda_t^U t) \\ &= \mathbb{P}_k(S_1 * S_2 * \dots * S_k(t) \geq \lambda_t^U t) \\ &\leq \inf_{u_1 \leq \dots \leq u_k} \mathbb{P}_k \left(\sum_{j=1}^k S_j(u_{j-1}, u_j) \geq \lambda_t^U t \right). \end{aligned} \quad (16)$$

For $u_1 \leq \dots \leq u_k$ we can expand the probabilities as

$$\begin{aligned} \mathbb{P}_k \left(\sum_{j=1}^k (u_j - u_{j-1} - V_j(u_{j-1}, u_j)) \geq \lambda_t^U t \right) \\ \leq E \left[\prod_{j=1}^k e^{-\theta V_j(u_{j-1}, u_j)} \right] e^{\theta(1-\lambda_t^U)t}, \end{aligned}$$

after using again the Chernoff bound.

At this point, we rearrange the terms in the product as in Eq. (15) and proceed as for the lower bound using Hölder's inequality first and then the independence property of V_j 's over non-overlapping intervals. We immediately get

$$\mathbb{P}_k(D(t) \geq \lambda_t^U t) \leq (b_{-\gamma\theta}^t)^{\frac{1}{\gamma}} e^{\theta(1-\lambda_t^U)t},$$

using the Laplace transform from Section IV.

Finally, for some fixed violation probability ε , the lower bound λ_t^L and the upper bound λ_t^U follow by the change of probability measure $\mathbb{P} = \sum_k \tilde{\pi}_k \mathbb{P}_k$, which completes the proof. \square

A. Sensitivity to N_j 's

Here, we discuss the capacity's sensitivity to the distribution of the number of neighbors N_j 's (herein referred to as N). We focus on the asymptotic capacity from Eq. (12), and it is thus sufficient to consider a single-hop transmission.

Let us firstly perform some preliminary calculations. Using Jensen's inequality we get that

$$\lambda = \sum_{l \geq 2} \pi_l p (1-p)^{l-1} \geq p (1-p)^{E[N]-1}.$$

The maximum in the last term is attained for $p = \frac{1}{E[N]}$, and thus $\lambda = \Omega(1/E[N])$. For the same value of p , and using an upper bound on Jensen's inequality (see Theorem 1.2 in Simic [31]), we get $\lambda = \mathcal{O}(1/E[N])$, and therefore the asymptotic capacity scales as

$$\lambda = \Theta \left(\frac{1}{E[N]} \right). \quad (17)$$

This scaling also holds in the case of a static network in which N is a constant, i.e., $\mathbb{P}(N \neq E[N]) = 0$.

An improved scaling law can be obtained by assuming that, on every sample-path ω , all the nodes are aware of their densities N_ω , and set their transmission probabilities to $p_\omega = \frac{1}{N_\omega}$; this probability is now a *random measure*. A lower bound on the asymptotic capacity is

$$\lambda = \sum_{l \geq 2} \pi_l \frac{1}{l} \left(1 - \frac{1}{l} \right)^{l-1} \geq \sum_{l \geq 2} \pi_l \frac{1}{le} = \frac{1}{e} E \left[\frac{1}{N} \right],$$

after using $(1 - \frac{1}{l})^{l-1} \geq \lim_{n \rightarrow \infty} (1 - \frac{1}{n})^{n-1} = \frac{1}{e}$. In turn, an upper bound is

$$\lambda = \sum_{l \geq 2} \pi_l \frac{1}{l} \left(1 - \frac{1}{l} \right)^{l-1} \leq \sum_{l \geq 2} \pi_l \frac{1}{l} = E \left[\frac{1}{N} \right].$$

Therefore, with the sample-path neighbor-aware probabilities assumption, the asymptotic capacity scales as

$$\lambda = \Theta \left(E \left[\frac{1}{N} \right] \right). \quad (18)$$

The same scaling is achieved by an ideal distributed scheduling mechanism, such as TDMA (which can be modelled as $X_i(t) = 1$ if $t \bmod N = l - 1$, and $X_i(t) = 0$ otherwise, in Eq. (5)) or by 802.11 DCF, since the transmission probabilities p implicitly scale as $\Theta(\frac{1}{N})$ (Bianchi [5]). Therefore, CSMA and TDMA networks achieve the same scaling, up to the assumptions discussed in Section III on CSMA; this result was previously shown to hold in the particular case of binomial node densities (i.e., a underlying uniform geometry) and an idealized CSMA model (see Chau *et al.* [10]).

Inspecting the two scalings from Eqs. (17) and (18), with Jensen's inequality $E[N]E[\frac{1}{N}] \geq 1$, reveals that the latter is asymptotically bigger. Therefore, the capacity in a random network with neighbor-aware transmission probabilities is asymptotically bigger than the capacity of a random network with a fixed value for p , or a static network with optimally adjusted p .

The discrepancy between the two scaling laws raises the interesting question on the gain-maximizing distribution of N which maximizes the capacity from Eq. (18). To prevent trivial scenarios, such as when there are only two nodes in the network ($\pi_2 = 1$), we look for the distribution of N relative to the normalized, or aligned, static scenario with a

Dist.	π_l	$E[N]$	$E\left[\frac{1}{N}\right]$	Gain
G-M	(20)	$\Theta(n)$	$\Theta(1)$	$\Theta(n)$
Unif.	$\frac{1}{n}$	$\Theta(n)$	$\Theta\left(\frac{\log n}{n}\right)$	$\Theta(\log n)$
Bnom.	$\binom{n}{l} \left(\frac{r}{q}\right)^l q^n$	$\Theta(n)$	$\Theta\left(\frac{1}{n}\right)$	$\Theta(1)$
Harm.	$\frac{\kappa}{l \log n}$	$\Theta\left(\frac{n}{\log n}\right)$	$\Theta\left(\frac{1}{\log n}\right)$	$\Theta\left(\frac{n}{\log^2 n}\right)$
Harm.	$\frac{\kappa / \log n}{(n-1)}$	$\Theta(n)$	$\Theta\left(\frac{1}{n}\right)$	$\Theta(1)$
Hv.tld.	$\frac{\kappa}{l^2}$	$\Theta(\log n)$	$\Theta(1)$	$\Theta(\log n)$
Hv.tld.	$\frac{\kappa}{(n-l)^2}$	$\Theta(n)$	$\Theta\left(\frac{1}{n}\right)$	$\Theta(1)$
Sbexp.	$\frac{\kappa}{l^3}$	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$
Sbexp.	$\frac{\kappa}{(n-l)^3}$	$\Theta(n)$	$\Theta\left(\frac{1}{n}\right)$	$\Theta(1)$
Geom.	κa^l	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$
Geom.	κa^{n-l}	$\Theta(n)$	$\Theta\left(\frac{1}{n}\right)$	$\Theta(1)$

TABLE I. CAPACITY RANDOMNESS GAINS FOR VARIOUS DISTRIBUTIONS $\pi = (\pi_1, \dots, \pi_n)$; n IS THE MAXIMUM VALUE OF N .

constant number $E[N]$ of nodes (note that the random and static scenarios are aligned in that the number of nodes are identical, on average). We are thus interested in maximizing the normalized randomness gain

$$\operatorname{argmax}_N \frac{\text{Eq. (18)}}{\text{Eq. (17)}} = \operatorname{argmax}_N E[N] E\left[\frac{1}{N}\right], \quad (19)$$

subject to the sample-path node-aware probabilities assumption. The next theorem provides the solution.

Theorem 2: (GAIN-MAXIMIZING DISTRIB. IN EQ. (19))

Denote $n = n_{\max}$ (the maximum value of N). Then Eq. (19) is maximized by the distribution

$$\pi_2 = \frac{n - E[N]}{n - 2}, \quad \pi_n = \frac{E[N] - 2}{n - 2}, \quad \pi_i = 0 \quad (i = 1, 3, \dots, n-1) \quad (20)$$

when both n and $E[N]$ are fixed.

For the proof see the Appendix. The intuition behind the bimodal distribution is that the increase rate in capacity by lowering the number of neighbors is larger than the decrease rate in capacity by increasing the number of neighbors. Note also that the distribution maximizes not only Eq. (19), but also the capacity when both n and $E[N]$ are fixed, under the sample-path node-aware probabilities assumption.

Table I illustrates the scaling of $E[N]E\left[\frac{1}{N}\right]$ from Eq. (19), for various distributions $\pi_l = \mathbb{P}(N = l)$. κ 's are normalization constants, $r = 1 - q$ for the binomial distribution (Bnom.), and $0 < a < 1$ for the geometric distribution (Geom.). There are two versions of harmonic (Harm.), heavy-tailed (Hv.tld.), subexponential (Sbexp.), and geometric distributions, depending on the hops' density; for instance, row four models a sparse situation with higher densities assigned to smaller number of nodes, whereas row five models the opposite situation.

The reported gain in the last column is relative to the expectation $E[N]$. Note that the gain-maximizing (G-M) distribution from Eq. (20) with $E[N] = \Theta(n)$ and the heavy-tailed distribution modelling sparse situations achieve the maximum relative gain, further supporting the observation after Theorem 2. The uniform (Unif.) distribution achieves the same gain $\Theta(\log n)$ as the first heavy-tailed distribution, but that is relative to an asymptotically larger expected number of nodes, i.e., $\Theta(n)$ vs. $\Theta(\log n)$.

Summarizing the results, we conclude that the asymptotic (in the time scale) per-flow capacity is fundamentally influ-

enced by randomness in network geometry, *but only* under the sample-path node-aware probabilities assumption. This assumption requires explicit overhead for both TDMA and Aloha schemes, and it is implicitly satisfied by 802.11 DCF since the nodes' transmission probabilities are $p_\omega = \Theta(1/N_\omega)$ in steady-state (Bianchi [5]). Similar conclusions can be drawn on the non-asymptotic capacity as well, since the bounds from Eqs. (10) and (11), with $\theta = \Theta(t^{-\zeta})$, deviate from Eq. (12) by constant terms depending on the time scale and/or the number of hops. Another interesting observation is that, according to Table I, the randomness gain is non-trivial as long as π_l does not decay faster than κ/l^2 ; an open problem concerns the general distribution of N for which the gain is $\omega(1)$.

Let us clarify the apparent contradiction between the above observation that "randomness increases the per-flow capacity" and the folk principle that "*determinism minimizes the queue*" from queueing theory (Humblet [21]) (which agrees in particular with the fact that randomness decreases the network capacity [18]). The reason is that our network model is slightly different than the one from [18], specifically by fixing both the source and the destination and letting the rest be random. Recall that our network model is deliberately tailored to directly study the per-flow rather than the network capacity.

B. Sensitivity to K

Here, we analyze the role of the distribution of the number of hops K on the lower bound of the non-asymptotic capacity from Eq. (10); as already pointed out, the other capacity results (i.e., the upper bound from Eq. (11) and the asymptotic one from Eq. (12)) are invariant to K . Despite this apparent incompleteness (we only analyze the lower bounds), we conjecture that the obtained scaling laws herein hold for upper bounds as well, given the previous observation that the lower bound captures the right scaling in K .

Analyzing the scaling law in K in Eq. (10) yields a trivial result, i.e., $\Theta(1)$, since the limit in t must be taken simultaneously (recall that $t \geq k_{\max}$ in Theorem 1). Informally, note that if $t = \Theta(E[K]^{1-\zeta})$, for $\zeta > 0$, then $\lambda_t^L = 0$ because there are insufficient time slots to carry packets from the source to the destination. To get more interesting results, let us properly scale $t = \Theta(E[K]^{1+\zeta})$ and $\gamma = \Theta(1)$, for some large $E[K]$. Then, the lower bound λ_t^L decays as

$$\lambda_t^L = \Omega(-\log E[K]) ,$$

in the case of a static network with a fixed number $E[K]$ of hops. In turn, in the case of a network with a random number of hops K , λ_t^L decays as

$$\lambda_t^L = \Omega(-E[\log K]) .$$

Jensen's inequality ($E[\log K] \leq \log E[K]$) implies that the lower bound on capacity in a random scenario is asymptotically bigger than in a static scenario. As in the previous subsection, this discrepancy raises the problem of the gain-maximizing distribution of K which maximizes the randomness gain, defined here as

$$\operatorname{argmax}_K (\log E[K] - E[\log K]) . \quad (21)$$

Before we provide the answer, in the next theorem, let us remark that the randomness gain is now defined in terms of

a *difference*, and not of a *ratio* as in Eq. (19). The reason stands in the contribution of K and N to the e2e capacity: the former has an additive effect (i.e., it affects the *cumulative throughput*), whereas the latter has a multiplicative effect (i.e., it affects the *throughput rate*).

Theorem 3: (GAIN-MAXIMIZING DISTRIB. IN EQ. (21)) Denote $k = k_{\max}$ (the maximum value of K). Then Eq. (21) is maximized by the distribution

$$\tilde{\pi}_1 = \frac{k - E[K]}{k - 1}, \quad \tilde{\pi}_k = \frac{E[K] - 1}{k - 1}, \quad \tilde{\pi}_i = 0 \quad (i \neq 1, k) \quad (22)$$

when both k and $E[K]$ are fixed.

The proof is similar to the proof of Theorem 2 and it is omitted.

Similar as in Theorem 2, the intuition for the bimodal distribution is that the rate at which capacity increases by lowering the number of hops is bigger than the rate at which capacity decreases by increasing the number of hops. Note that the distribution from Eq. (22) maximizes not only the randomness gain from Eq. (21) but also the capacity (its lower bound) when both k and $E[K]$ are fixed. Also, note that the randomness gain of the distribution from Eq. (22) can be asymptotically larger than a constant (e.g., $\Theta(\log \log k)$ vs. $\Theta\left(\frac{\log^2 k}{k}\right)$ when $E[K] = \Theta(\log k)$).

As far as other distributions are concerned, recall that in the case of the Uniform distribution for the number of neighbors, Table I reports a randomness gain of $\Theta(\log n)$. In contrast, there is no randomness gain in the case of the Uniform distribution for the number of hops. The reason is that while increasing the number of neighbors has a pronounced effect on capacity, increasing the number of hops has a much more moderate effect due to spatial reuse. This also indicates that spatial correlations are much less sensitive to randomness, as opposed to temporal correlations as observed in the previous subsection. As another closely related example, the first heavy-tailed distribution for N from Table I yields a $\Theta(\log n)$ gain in Eq. (19). In turn, using the convergence of the series $\sum_{l=1}^k \frac{\log l}{l^2}$, the same heavy-tailed distribution for K has the gain $\Theta(\log \log k)$ in Eq. (21).

The above example leads us to speculate that for specific distributions of N and K , the (per-flow) capacity gain due to randomness in the number of hops is the logarithm of the capacity gain due to randomness in nodes' densities.

C. Sensitivity to γ

Here, we briefly investigate the role of the dependency parameter γ on the non-asymptotic capacity.

Note firstly that taking a limit in γ would require taking limits in both $E[K]$ and t . As in the previous subsection, let $t = \Theta(E[K]^{1+\zeta})$ for some large $E[K]$, but take now $\gamma = \Theta(E[K])$ which models a high degree of dependencies amongst N_j 's (including the worst-case when all hops' densities are correlated to each other). Then, the lower bound decays as

$$\lambda_t^L = \Omega(-\log E[K]) \quad , \quad (23)$$

i.e., the logarithm is the price for assuming a high degree of spatial dependencies in the network. See also Burchard *et*

al. [7] where the same logarithmic factor is the additional decay on e2e delays (in a $\Theta(\cdot)$ sense) in networks with Markovian type of traffic, due to spatial dependencies; recall Item O4 from the Introduction.

Wrapping up, the main observations from this section are summarized in Items O1-O4 from the Introduction. These observations raise an interesting tradeoff between per-flow fairness/delay vs. capacity metrics, which has been addressed at the network level in the particular case of uniform geometries (see Grossglauser and Tse [17], Neely and Modiano [27], and Sharma *et al.* [30]). We believe that a further understanding of this tradeoff at the per-flow level, in networks with general geometries, can inspire distributed network algorithms emulating randomness, in order to provide differentiated per-flow services while maintaining a certain level of performance at the network level. More concisely, how could one leverage the randomness in the network geometry, given its advantageous impact on per-flow capacity?

VI. CONCLUSIONS

We have derived closed-form per-flow capacity results, in terms of both upper and lower (non-)asymptotic bounds, on a multi-hop path with a fixed source-destination pair. The key aspect of these results is that they apply to networks with *general random geometries* and various degrees of spatial correlations. By exploiting the simple analytical forms of the obtained results, we have quantified the beneficial impact of randomness in geometry on the per-flow capacity metric. In particular, we have shown that different distributions of hops' densities with normalized averages can lead to gaps in capacity scaling laws as large as $\Theta(n)$. We have identified a logarithmic detrimental factor of spatial correlations, and we have further observed a logarithmic relationship between spatial and temporal correlations.

Beyond the intuitively obvious message that randomness matters, this paper strives to communicate how does randomness *quantitatively* matter by analyzing a broad range of distributions. The collected observations jointly raise the awareness that the restriction to the widely-adopted uniform geometry model can be quite misleading. Moreover, these observations open the conceivably practical and yet challenging research problem of increasing per-flow capacity by leveraging the randomness in network geometry.

REFERENCES

- [1] N. Abramson. The Aloha system: another alternative for computer communications. In *Proceedings of AFIPS Joint Computer Conferences*, pages 281–285, 1970.
- [2] I. Akyildiz and X. Wang. *Wireless Mesh Networks*. Wiley, 2009.
- [3] G. Alfano, M. Garetto, E. Leonardi, and V. Martina. Capacity scaling of wireless networks with inhomogeneous node density: Lower bounds. *IEEE/ACM Transactions on Networking*, 18(5):1624–1636, Oct. 2010.
- [4] J. G. 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 Communications Magazine*, 46(12):94–101, Dec. 2008.
- [5] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, Mar. 2000.
- [6] J.-Y. Le Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.

- [7] A. Burchard, J. Liebeherr, and F. Ciucu. On superlinear scaling of network delays. *IEEE/ACM Transactions on Networking*, 19(4):1043–1056, Aug. 2011.
- [8] A. Bušić, J.-M. Fourneau, and N. Pekergin. Worst case analysis of batch arrivals with the increasing convex ordering. In A. Horváth and M. Telek, editors, *Formal Methods and Stochastic Models for Performance Evaluation*, volume 4054 of *Lecture Notes in Computer Science*, pages 196–210. Springer, 2006.
- [9] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [10] S. Chau, M. Chen, and S. Liew. Capacity of large scale CSMA wireless networks. In *ACM Mobicom*, 2009.
- [11] F. Ciucu. On the scaling of non-asymptotic capacity in multi-access networks with bursty traffic. In *IEEE International Symposium on Information Theory (ISIT)*, 2011.
- [12] F. Ciucu, O. Hohlfeld, and P. Hui. Non-asymptotic throughput and delay distributions in multi-hop wireless networks. In *Allerton Conference on Communications, Control and Computing*, 2010.
- [13] F. Ciucu, R. Khalili, Y. Jiang, L. Yang, and Y. Cui. Towards a system theoretic approach to wireless network capacity in finite time and space. In *IEEE Infocom*, 2014.
- [14] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *IEEE International Workshop on Quality of Service (IWQoS)*, pages 261–270, 2006.
- [15] M. Franceschetti, O. Dousse, D. N. C. Tse, and P. Thiran. Closing the gap in the capacity of wireless networks via percolation theory. *IEEE Transactions on Information Theory*, 53(3):1009–1018, Mar. 2007.
- [16] Y. Gao, D.-M. Chiu, and J. C. S. Lui. Determining the end-to-end throughput capacity in multi-hop networks: methodology and applications. In *ACM Sigmetrics/Performance*, pages 39–50, June 2006.
- [17] M. Grossglauser and D. N. C. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, Aug. 2002.
- [18] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, Mar. 2000.
- [19] B. Hajek. The proof of a folk theorem on queuing delay with applications to routing in networks. *Journal of the ACM*, 30(4):834–851, Oct. 1983.
- [20] J. Hoydis, M. Petrova, and P. Mähönen. The effects of topology on the local throughput of ad hoc networks. *Ad-Hoc and Sensor Wireless Networks*, 7(3-4):337–347, 2009.
- [21] P. A. Humblet. Determinism minimizes waiting time in queues. Technical report, MIT Laboratory for Information and Decision Systems, LIDS-P-1207, 1982.
- [22] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer, 2008.
- [23] L. Kleinrock and F. A. Tobagi. Packet switching in radio channels: Part I—carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Transactions on Communications*, 23(12):1400–1416, Dec. 1975.
- [24] S. Kulkarni and P. Viswanath. A deterministic approach to throughput scaling in wireless networks. *IEEE Transactions on Information Theory*, 50(6):1041 – 1049, June 2004.
- [25] D. C. Lee and J. N. Tsitsiklis. The worst bulk arrival process to a queue. Technical report, MIT Laboratory for Information and Decision Systems, LIDS-P-2116, 1992.
- [26] V. Martina, M. Garetto, and E. Leonardi. Capacity scaling of large wireless networks with heterogeneous clusters. *Performance Evaluation*, 67:1203–1218, Nov. 2010.
- [27] M. J. Neely and E. Modiano. Capacity and delay tradeoffs for ad hoc mobile networks. *IEEE Transactions on Information Theory*, 51(6):1917–1937, June 2005.
- [28] E. Peralvalov, R. S. Blum, and D. Safi. Capacity of clustered ad hoc networks: How large is “large”? *IEEE Transactions on Communications*, 54(9):1672–1681, Sept. 2006.
- [29] B. Rogozin. Some extremal problems in the theory of mass service. *Theory of Probability & Its Applications*, 11(1):144–151, 1966.
- [30] G. Sharma, R. Mazumdar, and N. B. Shroff. Delay and capacity tradeoffs in mobile ad hoc networks: A global perspective. *IEEE/ACM Transactions on Networking*, 15(5):981–992, Oct. 2007.
- [31] S. Simic. On an upper bound for Jensen’s inequality. *Inequalities in Pure and Applied Mathematics*, 10(2), 2009.
- [32] S. Toumpis. Capacity bounds for three classes of wireless networks: asymmetric, cluster, and hybrid. In *ACM MobiHoc*, pages 133–144, 2004.
- [33] W. Whitt. On approximations for queues, I: Extremal distributions. Technical report, Institute for Systems Research, ISR-TR-1989-37, 1989.
- [34] D. Wu and R. Negi. Effective capacity: A wireless link model for support of quality of service. *IEEE Transactions on Wireless Communication*, 2(4):630–643, July 2003.
- [35] M. Xie and M. Haenggi. Towards an end-to-end delay analysis of wireless multihop networks. *Ad Hoc Networks*, 7(5):849–861, July 2009.
- [36] F. Xue and P. R. Kumar. Scaling laws for ad-hoc wireless networks: An information theoretic approach. *Foundations and Trends in Networking*, 1(2):145–270, July 2006.
- [37] K. Zheng, F. Liu, L. Lei, C. Lin, and Y. Jiang. Stochastic performance analysis of a wireless finite-state Markov channel. *IEEE Transactions on Wireless Communications*, 12(2):782–793, Feb. 2013.
- [38] H. Al-Zubaidy, J. Liebeherr, and A. Burchard. A (min, \times) network calculus for multi-hop fading channels. In *IEEE Infocom*, 2013.

APPENDIX

PROOF OF THEOREM 2. According to Eq. (18) we have to solve a linear program with the objective function

$$\max_{\pi_i} \sum_{i \geq 2} \frac{1}{i} \pi_i$$

subject to the constraints

$$\sum_{i \geq 1} i \pi_i = m, \quad \sum_{i \geq 1} \pi_i = 1, \quad \pi_i \geq 0.$$

Recall that $\pi_1 = 0$ (there are at least two nodes). Select some π_i ’s satisfying the above constraints and assume that there exists $i \in \{3, \dots, n-1\}$ such that $\pi_i \neq 0$. The idea is to appropriately redistribute the entire value of π_i to the extremes π_2 and π_n , and then show that the new distribution increases the objective function.

Let $x = \frac{n-i}{n-2} \pi_i$ and consider the new distribution π'_i similar to π_i except for

$$\pi'_2 = \pi_2 + x, \quad \pi'_i = 0, \quad \pi'_n = \pi_n + \pi_i - x.$$

One can check that the redistribution of π_i preserves the mean m . It remains to prove that

$$\frac{1}{2} \pi_2 + \frac{1}{i} \pi_i + \frac{1}{n} \pi_n \leq \frac{1}{2} (\pi_2 + x) + \frac{1}{n} (\pi_n + \pi_i - x).$$

By rearranging terms the inequality reduces to $i \geq 2$ which is true.

One can repeat the above procedures for all i in the set $\{3, \dots, n-1\}$ satisfying $\pi_i \neq 0$, and reduce the initial linear program to

$$\max_{\pi_2, \pi_n} \left\{ \frac{1}{2} \pi_2 + \frac{1}{n} \pi_n \right\},$$

subject to the constraints

$$2\pi_2 + n\pi_n = m, \quad \pi_2 + \pi_n = 1, \quad \pi_2 \geq 0, \quad \pi_n \geq 0.$$

The solution is given by the choice from Eq. (20). \square