Opportunistic Walks on Random Geometric Networks and Their Application in Scalability Analysis

Ali Dabirmoghaddam J. J. Garcia-Luna-Aceves
Department of Computer Engineering
University of California, Santa Cruz
Santa Cruz, California 95064
Email: {alid, jj}@soe.ucsc.edu

***Abstract*—Opportunistic routing is studied as a representative example of location-aware greedy routing schemes. The routing process between an arbitrary source-destination pair is modeled as a directed random walk between the two ends in the underlying graph. The mean number of transmissions as well as the average multi-hop distance between arbitrary nodes are unified under a conceptual measure called expected length of the opportunistic walk in the induced network graph. We model this quantity as the mean time to absorption in a finite-state Markov chain. An explicit closed-form expression is presented to approximate the results and tight bounds are given. The accuracy of the results predicted by the analytical model is verified through simulation experiments. We also demonstrate an application of the foregoing model in defining proximity-based social models and identifying classes of social networks that are scalable.**

I. Introduction

Originated by Gupta and Kumar's seminal analysis [1], it came to be believed that wireless networks are not fundamentally scalable in size due to mutual interference, concurrent transmissions and increased accumulation of the relaying traffic load throughout the network. Although various communication models were examined in that work, one important aspect being neglected was a realistic interaction paradigm between nodes. In particular, it was assumed that sources and destinations are chosen uniformly and randomly within the network. This assumption disregards the natural drivers ruling the quality of social relationships in real-world networks of people. One such important element is the geographical dispersion between nodes. Numerous studies have shown that the physical distance plays a crucial role in initiating social interactions among people in both online and offline worlds [2]–[5].

Despite the tight dependency and extensive overlap between the social networks and their underlying communication networks, due to the involved complexities, the mainstream literature has only studied the performance of such networks separately. Examples of studies on communication networks neglecting the latent social relationships are [1], [6], [7]. In contrast, several interaction patterns and social paradigms [8]–[10] are independently studied while the restrictions imposed by realistic underlying communication networks are overlooked.

In this work, we present an analytical model that uses the geographical distance between nodes to capture the interplay between realistic communication algorithms and social relationships in composite networks. In the past few years, extensive research have been conducted on hop count statistics of wireless networks with geographic routing [11]–[13]. We revisit this problem and provide a Markov-chain formulation for it. For the routing algorithm, we study Opportunistic Routing [14] (OR) as a generic example for the broad class of location-aware greedy forwarding schemes. We model this routing algorithm as a directed random walk on the physical graph of the network. We call such a process an *opportunistic walk* and using that, we show how the notions of hop count (HC) and transmission count (TX) can be unified under a conceptual measure we call *length of the opportunistic walk*.

By means of this analysis, we demonstrate that the per-hop progress towards destination can be approximated by an iid process. Such a process is described by a probability distribution that incorporates the collective impact of all important physical and geometrical properties of the network, *e.g.*, link quality, node density, radio coverage, *etc*. The ultimate derivation, presented by Theorem 1, bounds the expected length of the walk for all nodes within a certain but arbitrary maximum physical distance in a form expressed by the characteristic function of the foregoing distribution.

For the social aspect, we use a power-law distribution on geographical distance to specify the frequency and quality of inter-node interactions. We finally demonstrate that how the combination of two models can be used to identify classes of composite networks that exhibit scalability. In short, the major contributions of this work can be summarized as follows:

- Presenting a fast, efficient and highly accurate method to calculate the expected hop count (EHC) and expected number of transmissions (ETX) between source-destination pairs at arbitrary distances under a greedy forwarding scheme.
- Using the geographical distance as the key ingredient to interrelate the concepts of communication and social networks under realistic settings.
- Identifying classes of proximity-based social networks that result in scalable structures.

The remainder is organized as follows. Section II provides a formal description of the problem. Section III describes a general framework for solving the problem using Markov chains and Section IV provides convenient tools for an efficient solution. Section V validates the analytical model using simulation results. Section VI discusses an application of the foregoing model in analysis of scalability in wireless social networks. Finally, Section VII concludes the paper.

II. Problem Definition and Assumptions

Consider a multi-hop wireless network, $\mathcal{N}(X;r)$, in which $X \subset \mathbb{R}^d$ is an arbitrary point process specifying nodes positions in d-dimensional space and $r \in \mathbb{R}^+$ their common radio range. In this network, $\forall x_i, x_j \in X$, we assume an undirected edge (link) between i and j wherever $\|x_i - x_j\| \leq r$, where $\|\cdot\|$ is the norm of choice on $\mathbb{R}^d$, *e.g.*, Euclidean norm. For simplicity of notation, we shall use i to refer to the node placed at $x_i \in X$. Since in most applications, the exact configuration of points in X may not be known a priori, we assume that x_i's are independent and identically distributed as per some common probability density function in $\mathbb{R}^d$. We call such a construction a *Random Geometric Network* (RGN) [15].

In a RGN, we shall assume that nodes communicate based on an OR paradigm. As such, no static route is constructed proactively between any source-destination pair. Instead, nodes dynamically defer the choice of next relaying hop upon the reception of the packet by one of the neighbors of the current relay which is geographically closer to the destination. The closest recipient neighbor is then selected to act as the next relaying hop and the process repeats until the packet is delivered to the destination. Definition 1 provides a more formal description of what we shall refer to as an *opportunistic walk*, hereinafter.

Definition 1. *Let $\mathcal{W}(s,t) = \langle s, \ldots, t\rangle$ be a walk on $\mathcal{N}(X;r)$. We say $\mathcal{W}(s,t)$ is an opportunistic walk from s to t and denote it with $s \overset{\text{OW}}{\rightsquigarrow} t$ if $\|x_u - x_t\| \geq \|x_v - x_t\|$ for all ordered pairs (u,v) on $\mathcal{W}(s,t)$.*

Definition 1 suggests that an opportunistic walk is indeed a directed random walk in which an intermediate node cannot be any farther from the destination than any of its precedents along the walk. In this sense, what OR does in a RGN is to find an opportunistic walk between the two ends of the route. The hop count along the opportunistic route can thus be seen as the length of the opportunistic walk in the underlying graph. Also, taking the effect of retransmissions into account, the latter can also reflect the required number of packet transmissions to reach the destination. This briefly explains the reason why we are interested in the expected length of opportunistic walks in this paper. The following proposition highlights a useful property of opportunistic walks.

Proposition 1. *If $\mathcal{W}(s,t) = \langle s, \ldots, t\rangle$ is an arbitrary opportunistic walk from s to t on $\mathcal{N}(X;r)$, then $\forall u \notin \{s,t\}$ on $\mathcal{W}(s,t)$, the walk segment $\mathcal{W}(u,t) = \langle u, \ldots, t\rangle$ is also an opportunistic walk on $\mathcal{N}(X;r)$.*

Proof: Follows recursively from Definition 1. ■

We introduce definitions required for a formal description of the problem. Notation $\mathcal{B}(x;r)$ denotes a ball centered at x with radius r.

Definition 2. *For any nodes s and t in $\mathcal{N}(X;r)$ with $x_s, x_t \in X$ and $D(s,t) \coloneqq \|x_s - x_t\|$ denoting their distance, $\Phi(s,t) \coloneqq \mathcal{B}(x_t; D(s,t))$ defines the relaying region for all $s \overset{\text{OW}}{\rightsquigarrow} t$.*

Definition 3. *For any nodes s and t in $\mathcal{N}(X;r)$, we define $\phi(s,t) \coloneqq X \cap \Phi(s,t)$ to be the set of potential relays for $s \overset{\text{OW}}{\rightsquigarrow} t$. In fact, we say u is a potential relay for $s \overset{\text{OW}}{\rightsquigarrow} t$ in $\mathcal{N}(X;r)$ if $x_u \in \phi(s,t)$.*

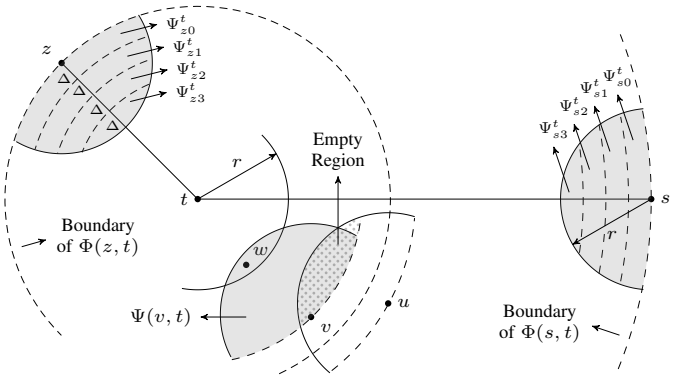

Fig. 1: An illustration of a RGN in 2-D space

Definition 4. *For any nodes s and t in $\mathcal{N}(X;r)$, we call $\Psi(s,t) \coloneqq \Phi(s,t) \cap \mathcal{B}(x_s;r)$ the hand-off region for s on all $s \overset{\text{OW}}{\rightsquigarrow} t$.*

Definition 5. *For any nodes s and t in $\mathcal{N}(X;r)$, a potential relay v in $\mathcal{N}(X;r)$ is a potential next hop for s on $s \overset{\text{OW}}{\rightsquigarrow} t$ if $x_v \in \psi(s,t) \coloneqq X \cap \Psi(s,t)$. We call $\psi(s,t)$ the set of potential next hops for s on $s \overset{\text{OW}}{\rightsquigarrow} t$.*

From Definitions 1 and 2, it is easy to see that for any node s and a fixed destination t, the relaying region on $s \overset{\text{OW}}{\rightsquigarrow} t$ shrinks as $\|x_s - x_t\| \to 0$. Similar implication comes from Definition 4 for the hand-off region. For better illustration, Fig. 1 provides a graphical display of a RGN in 2-D space where node t is the destination. The dashed lines specify the boundaries of relaying regions and the shaded areas illustrate the hand-off regions for corresponding relays. In this example, nodes v and w are potential relays for both $s \overset{\text{OW}}{\rightsquigarrow} t$ and $z \overset{\text{OW}}{\rightsquigarrow} t$, while node u is a potential relay only for $s \overset{\text{OW}}{\rightsquigarrow} t$.

III. The Markov Chain Formulation

The routing criterion in OR is to progressively reduce the remaining distance to the destination. Hence, at any intermediate stage and from the routing protocol's perspective, all nodes that are at the same distance from the destination are equally good to act as a potential next hop. In this sense, each stage of the walk can uniquely be described by a single measure, that is the remaining physical distance to the destination. Due to the greedy nature of OR, the remaining distance turns out to be a non-increasing quantity over time. To avoid confusion, we choose to use the complementary measure instead, *i.e.*, the total progress towards destination, to specify the states of the walk. By this convention, we shall obtain a more descriptive formulation.

According to Proposition 1, once a message comes at a certain distance to the destination along an opportunistic walk, the remaining section can be seen as another instance of opportunistic walk. Such intermediate stage is not, however, necessarily independent of the past history of the walk. For instance, consider the scenario depicted in Fig. 1, where node v is a potential next hop for u on $u \overset{\text{OW}}{\rightsquigarrow} t$. Assume that there is no other potential next hop for u in the dotted area. Part of the dotted area overlaps the hand-off region for v; therefore, v cannot have a potential next hop in the dotted subregion of its

hand-off region. Hence, the history from the previous stages *might* influence the future.

For simplicity of analysis, we shall neglect such interdependence between stages of the walk by characterizing the opportunistic walk as a memory-less stochastic process that exhibits Markov property. This assumption, of course, limits the accuracy of the results; however, we show later that this approximation imposes almost no adverse impact on networks with moderate to high density.

Due to its continuous nature, physical distance cannot directly be used to define the state space of a finite-state Markov chain. As a result, we first quantize the distance with a fixed step size of $\Delta \ll r$ to obtain the discretized progress towards destination. This treatment can be used in fact to describe a finite-state Markov chain. Indeed, if $D = \|x_s - x_t\|$ denotes the physical distance between source and destination, then $\mathcal{S}(s,t) = \{0, 1, \ldots, D_\Delta\}$ specifies the state space of the Markov chain, where $D_\Delta \coloneqq \lceil D/\Delta \rceil$ and $|\mathcal{S}| = m$. In such a chain, state $\sigma_u \in \mathcal{S}$ corresponds to the set of all potential relays u for which $x_u \in \mathcal{B}(x_t; \sigma_u\Delta) - \mathcal{B}(x_t; (\sigma_u + 1)\Delta)$ and $\psi(u,t) \neq \varnothing$. Note that the latter assumption is required to ensure the convergence of the expected length of the walk.

A. Specifying the State Transition Probability Distribution

Since packet forwarding is progressive in OR, a state $p \in \mathcal{S}$ can transit to any state $q \in \mathcal{S}$ with positive probability only if $0 \leq q - p \leq r_\Delta \coloneqq \lceil r/\Delta \rceil$. In order to quantify the transition probabilities for some potential relay u for $s \overset{\text{OW}}{\rightsquigarrow} t$ where $\|x_u - x_t\| \geq r$, we partition the hand-off region $\Psi(u,t)$ into $k = r_\Delta + 1$ mutually exclusive subregions $\{\Psi^t_{u0}, \Psi^t_{u1}, \ldots, \Psi^t_{u,k-1}\}$ (see Fig. 1) such that $\cup_{i=0}^{k-1} \Psi^t_{ui} = \Psi(u,t)$. If σ_u denotes the state of the Markov chain to which node u belongs, then a potential next hop in subregion $\Psi^t_{ui} \coloneqq \mathcal{B}(x_t; (\sigma_u + i + 1)\Delta) - \mathcal{B}(x_t; (\sigma_u + i)\Delta)$ corresponds to state $\sigma_u + i$ of the Markov chain, for any $0 \leq i < k$. The Markov chain transits from state σ_u to $\sigma_u + i$ if the following conditions are fulfilled by the end of next transmission trial:

- Some potential next hop for u in Ψ^t_{ui} has successfully received the packet from u; and
- for all $i < j < k$, either
  - there exists no potential next hop for u in Ψ^t_{uj}; or
  - no successful transmission is made to any potential next hop for u in Ψ^t_{uj}.

Define $\psi^t_{ui} := X \cap \Psi^t_{ui}$ to be the set of potential next hops for u in the i^{th} subregion, and let $\xi^t_{ui} \subseteq \psi^t_{ui}$ denote the set of potential next hops for u in Ψ^t_{ui} which successfully receive the packet from u. Then,

$$\begin{aligned}\Pr\left\{\xi^t_{ui} \neq \varnothing \,\middle|\, \psi(u,t) \neq \varnothing\right\} = \\ \Pr\left\{\exists \text{ node } v \text{ s.t. } x_v \in \psi^t_{ui} \,\middle|\, \psi(u,t) \neq \varnothing\right\} \times \\ \Pr\left\{\text{successful transmission from } u \text{ to } v\right\}. \quad (1)\end{aligned}$$

In the following, we provide some insights as to how the terms on the right-hand-side of Equation (1) can be evaluated.

1) The Existence of a Potential Next Hop: Depending on type of the underlying point process according to which nodes are distributed in the RGN, the probability of existence of a potential next hop in each subregion of $\Psi(u,t)$ can accordingly be quantified. Without loss of generality, in the following, we focus on a stationary Poisson point process with intensity ρ. In this case,

$$\Pr\left\{\psi^t_{ui} \neq \varnothing \,\middle|\, \psi(u,t) \neq \varnothing\right\} = \frac{1 - \exp\left(-\rho\left|\Psi^t_{ui}\right|\right)}{1 - \exp\left(-\rho\left|\Psi_{(u,t)}\right|\right)}, \quad (2)$$

where $|\Psi_{(\cdot)}|$ denotes the Lebesgue measure of the (sub)region.

The explicit evaluation of (sub)regions is not straightforward. Fig. 2a belongs to a 2-D example in which the areas of subregions are numerically calculated. As clearly seen, the probability of finding a potential next hop in farther subregions grows with the relay's distance to the destination (lower curves). In contrast, it becomes less likely to find a potential next hop in closer subregions as the relay goes farther from the destination (upper curves). Such a monotonic behavior indicates that it gradually becomes harder to find a potential next hop along an opportunistic walk closer to the destination. In other words, the walk's tendency to remain in its existing state increases when it comes closer to the destination. From Fig. 2a, it is perceived that for a fixed destination t, a relay u has the lowest chance to find a potential next hop in its farmost hand-off subregion when $\|x_u - x_t\| \to r^+$. As $\|x_u - x_t\| \to \infty$, however, this probability monotonically increases while being upper-bounded by a horizontal asymptote. We will refer to these limiting cases as r-distance and ∞-distance, respectively. Fig. 2b illustrates the r-distance and ∞-distance probabilities of finding a potential next hop in each of the 40 hand-off subregions for a relay with $r = 40$.

In the light of these observations, henceforth, we shall confine our attention to the behavior of the system exclusively for the two limiting cases. In the following, we briefly review a well-established model for assessing link quality.

2) The Packet Reception Model: Aside from the existence of a potential next hop v for a relay u, it is also important to see how reliably u is be able to pass the packet to v. The successful transmission from u to v relies on many physical and environmental factors. Contrary to the common assumption of idealized perfect-reception-within-range in wireless networks, it is widely known that the packet reception rate largely depends on the ambient noise and the distance to the transmitter. A realistic framework for modeling the probability of successful packet reception is presented by Zuniga and Krishnamachari [16] in which the collective impact of the important radio and environmental parameters are considered. The model suggests the following formula for the probability of packet reception over distance d:

$$f_D(d) = \left(1 - \frac{1}{2}\exp\left(-\frac{\gamma(d)}{1.28}\right)\right)^{8\ell}, \quad (3)$$

where ℓ is the size of the transmitted packet in bytes and $\gamma(d)$ is the SNR measured in dB at distance d from the transmitter which can be evaluated independently knowing the transmitting power, noise floor and path-loss exponent. Based on this model, Fig. 2c depicts the probabilities of successful reception for 40 packet transmissions (each bubble represents

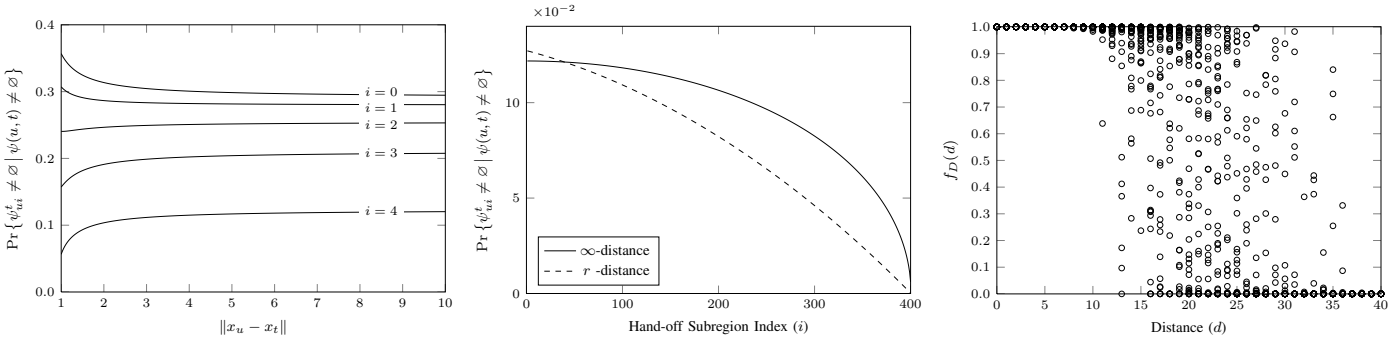

(a) Probability of finding potential next hop in subregions as a function of distance to destination

(b) Limiting probabilities of finding potential next hop in hand-off subregions

(c) Probability of packet reception with respect to distance from transmitter in presence of noise

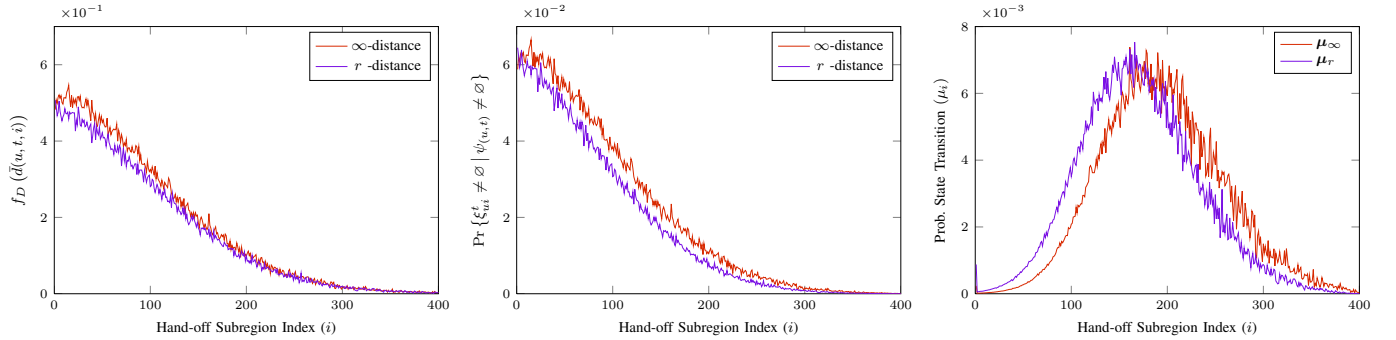

(d) The limiting probabilities of successful packet reception over hand-off subregions

(e) The limiting probabilities of having a potential next hop in a subregion receiving data

(f) The limiting distributions of transition into subsequent subregions

Fig. 2: The behavior of model components with respect to distance

an individual transmission at the given distance) where the distance between the transmitter and receiver changes from 0 to 40 m. The packets size is 400 bits and the noise standard deviation is 4.0. Denoting the average distance of points on Ψ_{ui}^t to x_u by $\bar{d}(u,t;i)$, Fig. 2d shows a Monte Carlo simulation of the limiting probabilities of successful transmission to imaginary potential next hops in 400 hand-off subregions.

Adopting the foregoing model for packet reception, Equation (1) can be approximated as follows:

$$\Pr\left\{\xi_{ui}^t \neq \varnothing \,\middle|\, \psi(u,t) \neq \varnothing\right\} \approx \Pr\left\{\psi_{ui}^t \neq \varnothing \,\middle|\, \psi(u,t) \neq \varnothing\right\} \times f_D\left(\bar{d}(u,t;i)\right) . \quad (4)$$

Fig. 2e clearly shows that this probability is higher for all subregions at ∞-distance than in r-distance.

B. The Limiting Transition Distributions

Putting all together, the probability of transition into state $\sigma_u + i$ from state σ_u, provided that relay u has at least one potential next hop, can be evaluated as follows:

$$\Pr\left\{S_{n+1} = \sigma_u + i \,\middle|\, S_n = \sigma_u,\, \psi_{(u,t)} \neq \varnothing\right\} = \left(1 - \Lambda_{(i,i)}\right) \Lambda_{(i+1,k-1)} + \mathbb{1}_{[i=0]}\, \Lambda_{(0,k-1)} , \quad (5)$$

where $\Lambda_{(j,k)} \coloneqq \prod_{i=j}^{k} \Pr\left\{\xi_{ui}^t = \varnothing \,\middle|\, \psi_{(u,t)} \neq \varnothing\right\}$. The last term on the right-hand-side of Equation (5) reflects the case where no potential next hop for u successfully receives the packet by the end of last transmission trial, and a retransmission has to be made from u. This contributes towards augmenting the probability of self-transition into state σ_u.

Note that the latter term could be neglected if EHC is the quantity of interest (instead of ETX). In fact, the difference between TX and HC arises from the impact of retransmissions. Such an effect can be omitted from our model assuming that by the end of each transmission trial, at least one potential next hop successfully receives the packet. Note that in such a case, $\Lambda_{(i,k-1)} = 0$. However, in order for Equation (5) to remain a valid probability distribution, Equation (4) needs to accordingly be adjusted to disregard the case when the given probability is set to zero for all hand-off subregions at the same time. In that case, the probability of that event would proportionally contribute towards the probability of other events, as follows:

$$\Pr\left\{S_{n+1} = \sigma_u + i \,\middle|\, \cdot\right\} = \frac{\left(1 - \Lambda_{(i,i)}\right) \times \Lambda_{(i+1,k-1)}}{\left(1 - \Lambda_{(0,k-1)}\right)} . \quad (6)$$

Computing the exact transition probabilities for all potential relays at distances from r to ∞ takes substantial effort. Fig. 2f illustrates the transition distributions for states corresponding to the two limiting distances (denoted by $\boldsymbol{\mu}_r$ and $\boldsymbol{\mu}_\infty$). As clearly evident from this figure, although the pace of approaching the destination state is faster with $\boldsymbol{\mu}_\infty$ than with $\boldsymbol{\mu}_r$, the difference between the two transition distributions is insignificant. With this remark in mind, we can approximate

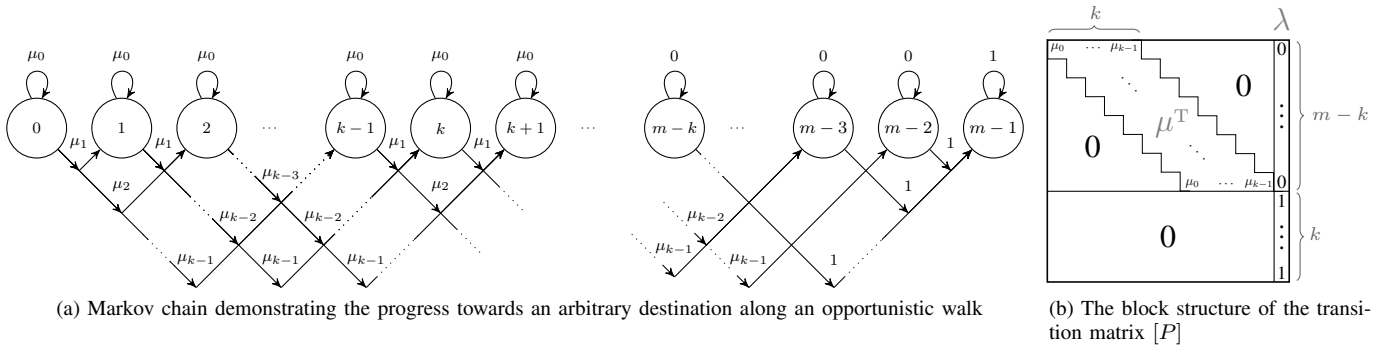

(a) Markov chain demonstrating the progress towards an arbitrary destination along an opportunistic walk

(b) The block structure of the transition matrix $[P]$

Fig. 3: The induced finite-state Markov chain and its corresponding transition matrix for an arbitrary random geometric network

the transition matrix of the Markov chain by assuming r-distance and ∞-distance transition distributions for all states on the chain. In the following subsection, we discuss this process in more details.

C. Constructing the Approximate Transition Matrix

Before getting to analyze the limiting transition scenarios, we first need to present a formal definition for a potential last relay on an opportunistic walk.

Definition 6. *We say a potential relay w for $s \overset{\text{OW}}{\rightsquigarrow} t$ is a potential last relay for $s \overset{\text{OW}}{\rightsquigarrow} t$ if $x_w \in \mathcal{B}(x_t; r)$.*

In Fig. 1, node w is a potential last relay for any node y on $y \overset{\text{OW}}{\rightsquigarrow} t$ where $\|x_y - x_t\| \geq \|x_w - x_t\|$. Naturally, a potential last relay w is able to reach destination t in finite time even if there is no additional potential relay y for which $\|x_y - x_t\| < \|x_w - x_t\|$. Among all potential relays for a certain opportunistic walk, only those that are not potential last relays (potential non-last relays) can have hand-off regions with exactly k subregions

In order to obtain an upper-bound and a lower-bound on the expected length of walk, we respectively use the r-distance and ∞-distance transition distributions for all states corresponding to non-last-relays, *i.e.*, the first $m-k$ states of the Markov chain. The transition distribution over the last k states, however, is entirely different due the fact that relays corresponding to those states have less than k subregions in their hand-off region. For simplicity of analysis, we assume that a potential last relay is always able to reach the destination using $O(1)$ transmissions. This simple assumption help eliminate some computational difficulties of the problem. Note that we can later relax this assumption by calculating the mean number of trials a potential last relay incurs (using Equation (3)), and adjusting the results accordingly. For the sake of brevity, however, we shall omit this part from our analysis.

With an identical transition distribution for the first $m-k$ states, if $i, j \in \mathcal{S}$ denote the states of two potential non-last-relays, then the probability of transition from i to j, denoted by $P_{i,j}$, depends only on $j-i$. In fact, we can write:

$$P_{i,j} = \begin{cases} \mu_{j-i} & 0 \leq j-i < k\,, \quad j \neq m\,, \\ 1 & m-k < i \leq m\,, \quad j = m\,, \\ 0 & \text{Otherwise}\,. \end{cases} \tag{7}$$

Using Equation (7), a finite-state Markov chain as depicted in Fig. 3a can be constructed for demonstrating the steps taken along an opportunistic walk on a RGN. The transition matrix, $[P]$, for such a Markov chain demonstrates a structure as depicted by Fig. 3b. From this figure, it is evident that the upper part of $[P]$ comprises a Toeplitz-like structure. As we shall shortly see in the following section, this particular structure exhibits many useful properties that can be used for simplifying our solution to the problem at hand.

IV. The Expected Length of Opportunistic Walks

The Markov chain illustrated in Fig. 3a contains m states among which state $m-1$ is absorbing and the rest are all transient states. The source and destination correspond to states 0 and $m-1$, respectively. It is clear that, for any finite value of m, the chain would eventually be trapped in state $m-1$ in finite time. The performance measures ETX and EHC can be regarded as the expected length of opportunistic walk between ending states (assuming Equations (5) and (6), respectively for transition distributions), which is, in fact, the mean time to absorption for the induced finite-state Markov chain. Letting $T_{m-1} := \inf\{t \geq 1 : X_t = m-1\}$ be a stopping time, the mean time to absorption can be computed recursively as:

$$\tau_0 = \mathbb{E}\big[T_{m-1} \,\big|\, X_0 = 0\big] = 1 + \sum_{j \in \mathcal{S}} P_{m-1,j}\, \tau_j \ . \tag{8}$$

Following similar logic for all other intermediate states, we can construct the following system of linear equations: $\boldsymbol{\tau} = \tilde{\boldsymbol{e}}_m + [P] \times \boldsymbol{\tau}$, in which $\tilde{\boldsymbol{e}}_m = (1, 1, \ldots, 1, 0)^{\mathrm{T}} \in \mathbb{R}^m$. Note that the last equation of this system has the vacuous form of $\boldsymbol{\tau}_{m-1} = 0 + \boldsymbol{\tau}_{m-1}$, with $\boldsymbol{\tau}_{m-1} = 0$. In order to solve this system, we let $[Q]$ be the same matrix as $[P]$ with its m^{th} row and column removed. Thus, we obtain:

$$\boldsymbol{\tau} = \Big([I] - [Q]\Big)^{-1} \times \boldsymbol{e}_{m-1}\,, \tag{9}$$

where $[I]$ is the identity matrix and $\boldsymbol{e}_{m-1} \in \mathbb{R}^{m-1}$ is a vector of all ones. Direct calculation of the given inverse matrix is inefficient, especially where $[Q]$ is large when it might become numerically singular. Nonetheless, we next present a series of useful insights that make these computations much faster and easier.

Lemma 1. *Let $[P]$ be an upper-triangular row-stochastic matrix as depicted in Fig. 3b, and $[Q]$ be the same matrix*

as $[P]$ with its last row and column removed. Then, matrix $[I]-[Q]$ is invertible, and $\big([I]-[Q]\big)^{-1} = \lim_{n\to\infty}\sum_{i=0}^{n}[Q^i]$.

Proof: $[Q]$ is upper-triangular and thus, all its eigenvalues are on the main diagonal. The elements on the main diagonal of $[Q]$ are inherited from the main diagonal of $[P]$. Thus, the eigenvalues of $[Q]$ are, in fact, the probabilities of self-transition into the transient states of the induced Markov chain and thus, are strictly less than one. Hence, $[I]-[Q]$ is upper triangular with nonzero elements along its main diagonal. This implies that $[I]-[Q]$ has a nonzero determinant and thus, is invertible. On the other hand, the spectral radius of $[Q]$ is less than one, guaranteeing the convergence $\lim_{n\to\infty}[Q^n] = [0]$. The rest of the proof easily follows by multiplying both sides of the given equation by $\big([I]-[Q]\big) \neq [0]$ and simple algebraic cancellations. ∎

Lemma 2. *The mean time to absorption for the finite-state Markov chain with transition matrix $[P]$ specified according to Fig. 3b is obtained as: $\boldsymbol{\tau} = \lim_{j\to\infty} j\cdot\boldsymbol{e}_m - \sum_{i=0}^{j-1}[P^i]_{:,m}$, where $[P^i]_{:,m}$ denotes the m^{th} (last) column of $[P^i]$.*

Proof: $[P]$ is a row-stochastic matrix and thus, $[P]\times\boldsymbol{e}_m = \boldsymbol{e}_m$. By induction on i, it is easy to observe that $[P^i]\times\boldsymbol{e}_m = \boldsymbol{e}_m$ for all $i\geq 0$. Therefore, powers of $[P]$ are also row-stochastic matrices. Using Lemma 1, we can rewrite Equation (9) in the following form: $\boldsymbol{\tau} = \sum_{i=0}^{\infty}[Q^i]\times\boldsymbol{e}_{m-1}$.

Note that $[Q^i]\times\boldsymbol{e}_{m-1}$ is a vector derived by summing up $[Q^i]$ along its rows. We know that $[Q]$ is formed by eliminating the last row and column from the upper-triangular row-stochastic matrix $[P]$ whose last row consists of only zero entries, except for $[P]_{m,m}$. Therefore, $[P^i]\equiv[Q^i]$ for all $i\geq 0$, where $\equiv$ signifies a one-to-one equivalence between corresponding elements of the two matrices. Thus, for every row $1\leq n<m$, $\sum_{q=1}^{m}[P^i]_{n,q} = \sum_{q=1}^{m-1}[Q^i]_{n,q} + [P^i]_{n,m} = 1$, which implies that $[Q^i]\times\boldsymbol{e}_{m-1} \equiv \boldsymbol{e}_m - [P^i]_{:,m}$. That is to say, the first $m-1$ elements from the two sides are equal. Also, since state $m-1$ is absorbing, $[P^i]_{m,m} = 1$, for all $i\geq 0$. This gives $\tau_{m-1} = 1 - [P^i]_{m,m} = 0$, which is obviously true. Consequently, it is obtained that $\boldsymbol{\tau} = \sum_{i=0}^{\infty}\big(\boldsymbol{e}_m - [P^i]_{:,m}\big) = \lim_{j\to\infty} j\cdot\boldsymbol{e}_m - \sum_{i=0}^{j-1}[P^i]_{:,m}$, which completes the proof. ∎

Using Lemma 2, finding the mean time to absorption in the induced Markov chain is reduced to calculating the sum of last columns of $[P^i]$. However, thanks to the special structure of $[P]$, we are able to further simplify the process of calculating these columns using the following lemma.

Lemma 3. *Given an upper-triangular row-stochastic matrix $[P]$ described according to Fig. 3b with vectors $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$, for every $i\geq 0$ we have:*

$$[P^{i+1}]_{n,m} = \begin{cases} \Big(\overleftarrow{\boldsymbol{\mu}}^{\star i}\star\boldsymbol{\lambda}_{\{\vartheta_i;1\}}\Big)_{\vartheta_i+n} & 1\leq n\leq m-k\,,\\ 1 & m-k<n\leq m\,,\end{cases}$$

where $\overleftarrow{\boldsymbol{\mu}}^{\star i}$ is the flipped version of the i^{th}-fold self-convolution of $\boldsymbol{\mu}$, and $\boldsymbol{\lambda}_{\{\vartheta_i;1\}}$ is vector $\boldsymbol{\lambda}$ right-padded with $\vartheta_i \coloneqq i(k-1)$ one-entries to comply with the length of $\overleftarrow{\boldsymbol{\mu}}^{\star i}$.

Proof: We consider two cases separately.

Case 1: $1\leq n\leq m-k$: The proof is by induction on i. If $i=0$, then $[P]_{n,m} = \big(\boldsymbol{\mu}^{\star 0}\star\boldsymbol{\lambda}_{\{0;1\}}\big)_{0+n} = \boldsymbol{\lambda}_n$, which is true. Also, if $i=1$:

$$\begin{aligned}[P^2]_{n,m} &= \textstyle\sum_{q=1}^{m}[P]_{n,q}\cdot[P]_{q,m}\\ &= \textstyle\sum_{q=1}^{m}\mu_{q-n}\cdot\lambda_q\\ &= \textstyle\sum_{q=1}^{m}\overleftarrow{\mu}_{n-q}\cdot\lambda_q\\ &= \textstyle\sum_{q=1}^{m}\overleftarrow{\mu}_{n-q}\cdot\big(\lambda_{\{\vartheta_1;1\}}\big)_{\vartheta_1+q}\\ &= \big(\overleftarrow{\boldsymbol{\mu}}\star\boldsymbol{\lambda}_{\{\vartheta_1;1\}}\big)_{\vartheta_1+n}\,.\end{aligned}$$

Let us assume that the given statement holds for $i=j\geq 1$. Then, for $i=j+1$ we have:

$$\begin{aligned}[P^{j+1}]_{n,m} &= \textstyle\sum_{q=1}^{m}[P]_{n,q}\cdot[P^k]_{q,m}\\ &= \textstyle\sum_{q=1}^{m}\mu_{q-n}\cdot\Big(\overleftarrow{\boldsymbol{\mu}}^{\star(j-1)}\star\boldsymbol{\lambda}_{\{\vartheta_{j-1};1\}}\Big)_{\vartheta_{j-1}+q}\\ &= \Big(\overleftarrow{\boldsymbol{\mu}}\star\big(\overleftarrow{\boldsymbol{\mu}}^{\star(j-1)}\star\boldsymbol{\lambda}_{\{\vartheta_{j-1};1\}}\big)\Big)_{\vartheta_j+n}\\ &= \Big(\overleftarrow{\boldsymbol{\mu}}^{\star j}\star\boldsymbol{\lambda}_{\{\vartheta_j;1\}}\Big)_{\vartheta_j+n}\,.\end{aligned}$$

Note that $\overleftarrow{\boldsymbol{\mu}}^{\star j}$ can be arbitrarily longer than $\boldsymbol{\lambda}$. This is why $\boldsymbol{\lambda}$ has to be one-padded at the end at least to the length of $\overleftarrow{\boldsymbol{\mu}}^{\star j}$ to ensure that no information is lost.

Case 2: $m-k<n\leq m$: The proof easily follows noting the fact that $[P]$ comprises only zero entries on its lower part; except for the ones along its last column. Thus, the lower part of $[P^i]$ remains unchanged for all $i>0$. ∎

There are some subtleties involved in selecting appropriate vector indices in Lemma 3. Such complications arise from the fact that convolution does not preserve the length of input vectors. In particular, the convolution $\Big(\overleftarrow{\boldsymbol{\mu}}^{\star i}\star\boldsymbol{\lambda}_{\{\vartheta_i;1\}}\Big)$ produces a vector in which indices $1+i(k-1)$ through $m-k+i(k-1)$ correspond to the first $m-k$ elements on $[P^i]_{:,m}$.

According to Lemma 3, for all powers of $i>0$, the last k elements on $[P^i]_{:,m}$ are always one. Consequently, by Lemma 2, it is entailed that the expected length of an opportunistic walk is always 1 from a potential last hop to destination and 0 from destination to itself, which are intuitive based on our assumptions. In the following, we shall exclude these quantities from our computations, merely focusing on the expected length of walks from potential non-last relays to destination, *i.e.*, the mean times to absorption for the first $\eta \coloneqq m-k$ states which we denote with vector $\boldsymbol{\tau}^{\dagger}$, henceforth.

Theorem 1. *The expected length of opportunistic walks from potential non-last relays to an arbitrary but fixed destination in a* RGN *is bounded by the mean time to absorption for the induced finite-state Markov chain with limiting transition probability vector $\boldsymbol{\mu}$, that can be computed as:*

$$\boldsymbol{\tau}^{\dagger}_{(\boldsymbol{\mu})} = \lim_{j\to\infty} j\cdot\boldsymbol{e}_{\eta} - \mathscr{F}^{-1}_{1:\eta}\left(\boldsymbol{\chi}_j\right)\,. \tag{10}$$

where,

$$\boldsymbol{\chi}_j[i] \coloneqq \begin{cases} (j-1)\cdot\mathscr{F}_{[1]}(\boldsymbol{\lambda}_{\{\bar{\vartheta}_j;1\}}) & i=1\,,\\ \left(\dfrac{1-\mathscr{F}^{j-1}_{[i]}\big(\boldsymbol{\mu}_{\{\bar{\vartheta}_j;0\}}\big)}{1\;-\;\mathscr{F}_{[i]}\big(\boldsymbol{\mu}_{\{\bar{\vartheta}_j;0\}}\big)}\right)^{*}\cdot\mathscr{F}_{[i]}\big(\boldsymbol{\lambda}_{\{\bar{\vartheta}_j;1\}}\big) & i>1\,.\end{cases} \tag{11}$$

Proof: Using Lemmas 2 and 3, we are able to express the mean time to absorption for the potential non-last-relays as an infinite series of serial convolutions. However, evaluating such a series is somewhat tricky, because different ranges of indices from each convolution term are useful in our computations. Nonetheless, with sufficient padding, we are able to appropriately align the variable-sized vectors so as to be able to meaningfully add them together, and then select the desired range from the sum at the end. For this, we first note that the desired range of indices for the i^{th} convolution term starts at ϑ_i; a quantity that grows with i. Also, the i^{th} term is longer in length than all the previous ones. Therefore, if we want to choose the appropriate range of indices after performing the summation, we must right-align all convolution terms before adding them together. Alternatively, instead of using a flipped version of $\boldsymbol{\mu}$ in our computations, we are able to use the original version of $\boldsymbol{\mu}$ with sufficient zero-padding on the right-side in self-convolutions and then do a regular summation on them. The resulting vector, however, would be the flipped version of what we are looking for. To obtain appropriately aligned vectors, we make sure to zero-pad each term to the length of the longest convolution term[1]. Therefore,

$$\overleftarrow{\boldsymbol{\tau}}^{\dagger}_{(\boldsymbol{\mu})} = \lim_{j\to\infty} j\cdot \boldsymbol{e}_{\eta} - \left(\left(\sum_{i=1}^{j-2} \boldsymbol{\mu}^{\star i}_{\{\bar{\vartheta}_j;0\}} \right) \star \boldsymbol{\lambda}_{\{\bar{\vartheta}_j;1\}} \right)_{\vartheta_j+1:\vartheta_j+\eta} \tag{12}$$

Using a discrete Fourier transform, we can convert the big convolution term inside the brackets into product to further simplify Equation (12). From harmonic analysis, note that the time-reversal of a given vector produces the conjugate of the corresponding version in the frequency domain. With this remark in mind, we construct vector $\boldsymbol{\chi}_j$ as follows:

$$\boldsymbol{\chi}_{j\,[\ell]} \coloneqq \left(\sum_{i=1}^{j-2} \mathscr{F}^{i}_{[\ell]} \left(\boldsymbol{\mu}_{\{\bar{\vartheta}_j;0\}} \right) \right)^{*} \times \mathscr{F}_{[\ell]} \left(\boldsymbol{\lambda}_{\{\bar{\vartheta}_j;1\}} \right) , \tag{13}$$

whereby we shall be able to have vector $\overleftarrow{\boldsymbol{\tau}}^{\dagger}$ backflipped to its right order after taking the inverse transform later on. Since $\boldsymbol{\mu}$ is a probability vector, $\mu_i \leq 1$ for all $1 \leq i \leq k$ and $\sum_i \mu_i = 1$. Also, we have $\left|\mathscr{F}_{[1]}(\boldsymbol{\mu})\right| = 1$ and $\left|\mathscr{F}_{[i]}(\boldsymbol{\mu})\right| < 1$ for all $i > 1$. Therefore, Equation (13) is a geometric series whose elements converge to the form given by Equation (11). From Equations (12) and (13) together, we obtain:

$$\mathscr{F}(\boldsymbol{\tau}^{\dagger}_{(\boldsymbol{\mu})}) = \lim_{j\to\infty} \mathscr{F}\left(j\cdot \boldsymbol{e}_{\eta}\right) - \boldsymbol{\chi}_j \ . \tag{14}$$

Finally, taking the inverse transform yields Equation (10). ■

The transition probability vector $\boldsymbol{\mu}$ is, in fact, a mass function. In statistics, the Fourier transform of a PMF is often referred to as *characteristic function* which exhibits many useful properties. The power of Theorem 1 lies in its ability to provide bounds on the expected length of opportunistic walks using only the characteristic function of the limiting transition distributions, *i.e.*, $\boldsymbol{\mu}_r$ and $\boldsymbol{\mu}_\infty$.

V. Simulation Results and Validation

We present simulation experiments to verify the correctness of our proposed model for the expected length of opportunistic walks in RGNs. To this end, a version of OR [14] is implemented in OMNeT++ [17], the discrete-event simulator. To ensure closeness of simulation results to reality, we have borrowed the simulation configuration from empirical measurements on MICA2 motes reported in [16], [18]. Using this setting, nodes are able to cover a radius of up to 40 meters with their radio. Nonetheless, for a lossy channel, as demonstrated earlier in Fig. 2c, the chance of successful packet reception is practically zero farther than 25 meters away from the transmitter.

[1] We denote $\boldsymbol{\mu}_{\{\bar{\vartheta}_j;0\}}$ when vector $\boldsymbol{\mu}$ is zero-padded to a final length of ϑ_j.

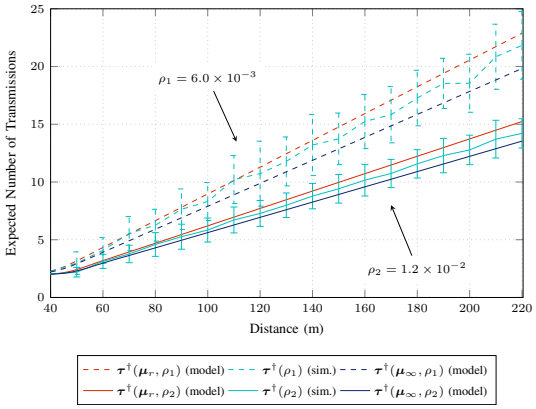

Fig. 4: Bounds on ETX (model vs. simulation)

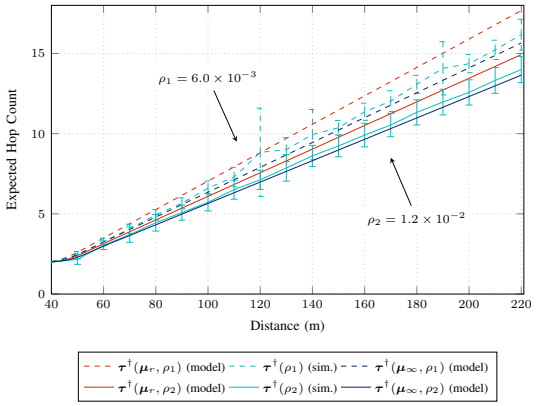

Fig. 5: Bounds on EHC (model vs. simulation)

Fig. 4, depicts ETX as a function of distance between source and destination, for two different node densities. The upper and lower lines show the analytical bounds computed from Theorem 1 in MATLAB™, whereas the middle curves show the simulation results in OMNeT++. In the simulation setup, each experiment includes transfer of 1000 packets between an arbitrary source-destination pair located at a certain distance from each other varying from 40 m to 220 m at steps of 10 m. For a higher fidelity, each simulation is run 100 times with different random seeds and the average is reported. Also, the error bars demonstrate the standard deviation for each simulation experiment. As clearly seen from Fig. 4, there is a perfect match between the simulation results and the analytical model. For sparser networks, however, although the simulation results demonstrate larger variations, the overall trend is accurately captured by the proposed model.

For an analysis of EHC, we simply ignore the effect of

retransmissions due to poor channel quality. A reevaluation of the limiting transition probability vectors $\boldsymbol{\mu}_r$ and $\boldsymbol{\mu}_\infty$ according to Equation (6) is needed to suppress the effect of self-transitions due to packet loss. Fig. 5 depicts the expected hop distance of nodes at physical distances given along the x-axis. Similar to the previous experiment, the model again exhibits a promising ability to capture the precise behavior of HC in all scenarios.

As previously discussed, Theorem 1 expresses the expected length of the walk as a limit to infinity. The limit variable (j) indeed determines how many powers of the transition matrix $[P]$ are taken into account in order for approximating the output vector $\boldsymbol{\tau}$. In our simulations, we verified that setting $j = 40$ is sufficient to guarantee a relative error of less than 0.01% in most cases.

VI. Scalability Analysis of Proximity-based Social Networks

In this section, we briefly discuss a simple application of the proposed model in examining the scalability of networks. Before we proceed, let us present a formal description of the proximity-based social networks.

Definition 7. *A proximity-based social network is a* RGN $\mathcal{N}(X;r)$ *in which all nodes* u *and* v *with* $x_u, x_v \in X$ *communicate with probability proportional to* $D(u,v)^{-\alpha}$ *for some* $\alpha \geq 0$*, where* $D(u,v) = \|x_u - x_v\|$.

In Definition 7, we adopt a simple yet realistic model for the frequency of pairwise interactions in RGNs. The inverse power model was first presented and verified by Latané *et al.* [2] in 1995 and has been widely used in many analyses ever since, *e.g.*, [9]. The parameter $\alpha \geq 0$ is the clustering exponent that controls how restrictive the model is in limiting the social relationships within a proximal neighborhood of the source. For example, when $\alpha = 0$, regardless of their physical distance, all nodes are equally likely to communicate (uniform communication model). On the other hand, as α grows, nodes demonstrate higher tendency to communicate with their proximate neighbors more often than with those at the far end of the network.

In order to obtain a distribution for the probability of having a social interaction at a given physical distance, we continue with our 2-D network. In that case and with the Poisson approximation, we note that the number of potential neighbors at a certain physical distance linearly grows with distance. Denoting the quantized distance $\delta_{[n]} = n\Delta$ for $n \geq 0$, we can obtain the following approximation for the probability mass function in our proximity-based social model.

$$\boldsymbol{f}_{\alpha[n]} \approx \frac{\delta_{[n]}^{1-\alpha}}{\sum_{t=0}^{D/\Delta} \delta_{[t]}^{1-\alpha}} = \frac{n^{1-\alpha}}{\sum_{t=0}^{D/\Delta} t^{1-\alpha}}\,, \tag{15}$$

where $D \coloneqq \max\big(D(u,v)\big)\,\forall u, v$ is the network diameter.

One interesting aspect to investigate, given a proximity-based social model, is the average number of hops (or transmissions) it takes for a typical packet to reach its destination. We refer to such quantity as the expected social path length which is defined as follows.

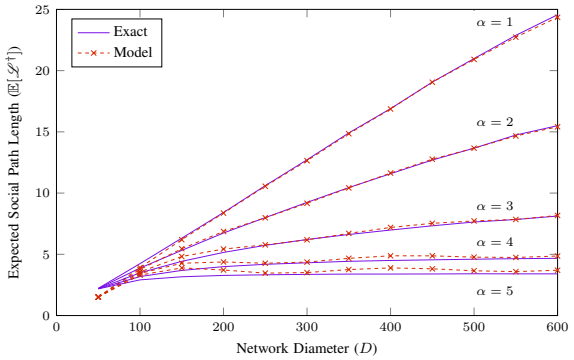

Fig. 6: ESPL as a function of the network diameter

Definition 8. *The expected social path length (*ESPL*) is the expected length of the opportunistic walk to any other node in a proximity-based social network identified with* $\boldsymbol{f}_\alpha$*, that is calculated as:* $\mathbb{E}[\mathscr{L}] = \sum_{n=0}^{D/\Delta} \boldsymbol{\tau}_{[n]} \times \overleftarrow{\boldsymbol{f}}_{\alpha[n]}$.

Excluding the one-hop neighbors from our analysis, we can take advantage of Theorem 1 to derive a more detailed expression for ESPL. Defining $\boldsymbol{f}_\alpha^\dagger \coloneqq \overleftarrow{\boldsymbol{f}}_{\alpha\{\vartheta_j;0\}}$ for $\delta_{[n]} > r$, we can write:

$$\begin{aligned}
\mathbb{E}[\mathscr{L}^\dagger] &= \lim_{j\to\infty} \sum_{i=1}^{\vartheta_j} \Big(j \cdot \boldsymbol{f}_{\alpha[i]}^\dagger - \mathscr{F}_{[i]}^{-1}(\chi_j) \cdot \boldsymbol{f}_{\alpha[i]}^\dagger \Big) \\
&= \lim_{j\to\infty} j - \sum_{i=1}^{\vartheta_j} \mathscr{F}_{[i]}^{-1}\big(\boldsymbol{f}_\alpha^\dagger\big) \cdot \chi_{j\,[i]} \\
&= \frac{D}{r} - \lim_{j\to\infty} \sum_{i=2}^{\vartheta_j} \mathscr{F}_{[i]}^{-1}\big(\boldsymbol{f}_\alpha^\dagger\big) \cdot \chi_{j\,[i]}
\end{aligned} \tag{16}$$

In the following, we investigate to see how different values of α affect the growth of ESPL with respect to the network size. For that purpose, while maintaining a fixed node density, we expand the network size by adding more nodes and calculate ESPL when different values of α are adopted in our social model.

Fig. 6 illustrates ESPL as a function of the network diameter when node density is 6.0×10^{-3}. The solid curves represent the exact values of ESPL derived from Definition 8. The dashed lines, however, are calculated from the proposed model in Equation (16). Even though for this analysis only the first five harmonic components have been used in the summation (*i.e.*, $\vartheta_j = 5$), still, as clearly seen, there is a good match between the model and the exact results.

From Fig. 6, it can be observed that as the network expands in size, ESPL also increases accordingly. However, the growth rate of ESPL diminishes as α increases. Essentially, ESPL reflects the average amount of network resources utilized for multi-hop packet forwarding and thus, is a useful measure in the analysis of network performance. Evidently, ESPL is a non-decreasing function of the network size; nevertheless, the network cannot obviously sustain a continuously increasing load forever as more nodes join in. Hence, we present the

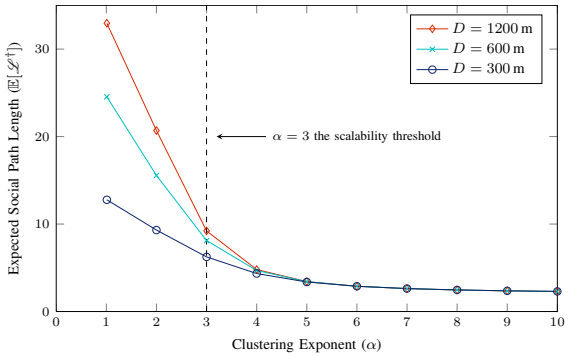

Fig. 7: ESPL as a function of the clustering exponent

following definition for social networks that can appropriately scale without significant loss in their performance.

Definition 9. *We say a proximity-based social network is scalable if* $\mathbb{E}[\mathscr{L}] < \infty$ *when* $D \to \infty$.

Based on Definition 9, a scalable network is a type of RGN in which the network, on average, incurs a finite number of transmissions (or hops) for information dissemination, no matter how large the network would grow.

A natural question arising here is for what values of α a proximity-based social network would be scalable. To answer this question, we note the growth of the expected length of opportunistic walks with the network diameter from Figg. 4 and 5. Using basic regression techniques, it is easy to verify the linear relationship between these quantities, at least for the typical size of networks used in nowadays applications. Approximating $\boldsymbol{\tau}_{[n]} \approx \beta\, \delta_{[n]}$ for some constant β independent of α and n, we have

$$\mathbb{E}[\mathscr{L}] \approx \lim_{D\to\infty} \sum_{n=0}^{D/\Delta} \beta\, \Delta \frac{n^{2-\alpha}}{\sum_{t=0}^{D/\Delta} t^{1-\alpha}}\,. \tag{17}$$

which converges to the form

$$\mathbb{E}[\mathscr{L}] \approx \beta\, \Delta\, \frac{\zeta(\alpha-2)}{\zeta(\alpha-1)}\,, \tag{18}$$

for $\alpha > 3$, where $\zeta(\cdot)$ is the Riemann zeta function.

To better understand this result, Fig. 7 shows the convergence speed of ESPL to the limiting value of 2 hops as the clustering exponent increases. As stated earlier, note that we only consider multi-hop communication in this analysis and hence, the limiting value 2. According to Fig. 7, in social networks with $\alpha \geq 5$, regardless of the network size, nodes barely get to communicate beyond their close neighborhood of two hops. For $\alpha = 3$, it is interesting to graphically observe that while ESPL grows with the network size, the growth rate is so small that even with an exponentially increasing network size, the differences in average social path lengths are insignificant.

VII. Conclusions and Outlook

A novel framework was presented for modeling of the mean hop distance and average transmission count along geographic routes in wireless multi-hop networks. By introducing the abstract concept of opportunistic walks, we were able to unify notions of hop distance and packet transmission count. We derived explicit tight bounds on the expected length of such walks and used these results to identify classes of proximity-based social networks that are scalable. The outlook for the future includes extended applications of the proposed model in performance and throughput analyses of wireless networks.

Acknowledgments

The authors gratefully thank the anonymous reviewers for their constructive comments on an earlier version of this manuscript. This research was sponsored in part by the Jack Baskin Chair of Computer Engineering at UCSC and by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053.

References

[1] P. Gupta and P. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.

[2] B. Latané, J. Liu, A. Nowak, M. Bonevento, and L. Zheng, "Distance Matters: Physical Space and Social Impact," *Personality and Social Psychology Bulletin*, vol. 21, no. 8, pp. 795–805, Aug. 1995.

[3] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins, "Geographic Routing in Social Networks," *PNAS*, vol. 102, no. 33, pp. 11 623–11 628, Aug. 2005.

[4] P. Fraigniaud and G. Giakkoupis, "On the Searchability of Small-World Networks with Arbitrary Underlying Structure," in *Proc. ACM STOC*, Cambridge, MA, Jun. 2010, pp. 389–398.

[5] L. Backstrom, E. Sun, and C. Marlow, "Find Me If You Can: Improving Geographical Prediction with Social and Spatial Proximity," in *Proc. WWW*, Raleigh, NC, Apr. 2010, pp. 61–70.

[6] M. Grossglauser and D. Tse, "Mobility Increases the Capacity of Ad Hoc Wireless Networks," *IEEE/ACM Trans. Netw.*, vol. 10, no. 4, pp. 477–486, Aug. 2002.

[7] S. Kulkarni and P. Viswanath, "A Deterministic Approach to Throughput Scaling in Wireless Networks," *IEEE Trans. Inf. Theory*, vol. 50, no. 6, pp. 1041–1049, Jun. 2004.

[8] D. Watts and S. Strogatz, "Collective Dynamics of 'Small-World' Networks," *Nature*, vol. 393, no. 6684, pp. 440–442, Jun. 1998.

[9] J. Kleinberg, "The Small-World Phenomenon: An Algorithm Perspective," in *Proc. ACM STOC*, Portland, OR, May 2000, pp. 163–170.

[10] M. Dietzfelbinger and P. Woelfel, "Tight Lower Bounds for Greedy Routing in Uniform Small World Rings," in *Proc. ACM STOC*, Maryland, DC, May 2009, pp. 591–600.

[11] S. Dulman, M. Rossi, P. Havinga, and M. Zorzi, "On the Hop Count Statistics for Randomly Deployed Wireless Sensor Networks," *IJSNet*, vol. 1, no. 1/2, pp. 89–102, Sep. 2006.

[12] S. Vural and E. Ekici, "On Multihop Distances in Wireless Sensor Networks with Random Node Locations," *IEEE Trans. Mobile Comput.*, vol. 9, no. 4, pp. 540–552, Apr. 2010.

[13] Z. Zhang, G. Mao, and B. Anderson, "On the Hop Count Statistics in Wireless Multi-hop Networks Subject to Fading," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 7, pp. 1275–1287, Jul. 2011.

[14] S. Biswas and R. Morris, "Opportunistic Routing in Multi-hop Wireless Networks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, pp. 69–74, Jan. 2004.

[15] M. Penrose, *Random Geometric Graphs*. New York, NY: Oxford University Press Oxford, 2003.

[16] M. Zuniga and B. Krishnamachari, "Analyzing the Transitional Region in Low Power Wireless Links," in *Proc. IEEE SECON*, Santa Clara, CA, Oct. 2004, pp. 517–526.

[17] A. Varga, "The OMNeT++ Discrete Event Simulation System," in *Proc. ESM*, Prague, Czech Republic, Jun. 2001.

[18] D. Faria, "Modeling Signal Attenuation in IEEE 802.11 Wireless LANs - Vol. 1," Stanford University, Tech. Rep. TR-KP06-0118, Jul. 2005.