

Efficient Routing in Intermittently Connected Mobile Networks: The Single-Copy Case

Thrasylvoulos Spyropoulos, *Student Member, IEEE*, Konstantinos Psounis, *Member, IEEE*, and Cauligi S. Raghavendra, *Fellow, IEEE*

Abstract—Intermittently connected mobile networks are wireless networks where most of the time there does not exist a complete path from the source to the destination. There are many real networks that follow this model, for example, wildlife tracking sensor networks, military networks, vehicular ad hoc networks (VANETs), etc. In this context, conventional routing schemes would fail, because they try to establish complete end-to-end paths, before any data is sent.

To deal with such networks researchers have suggested to use flooding-based routing schemes. While flooding-based schemes have a high probability of delivery, they waste a lot of energy and suffer from severe contention which can significantly degrade their performance. With this in mind, we look into a number of “single-copy” routing schemes that use only one copy per message, and hence significantly reduce the resource requirements of flooding-based algorithms. We perform a detailed exploration of the single-copy routing space in order to identify efficient single-copy solutions that (i) can be employed when low resource usage is critical, and (ii) can help improve the design of general routing schemes that use multiple copies. We also propose a theoretical framework that we use to analyze the performance of all single-copy schemes presented, and to derive upper and lower bounds on the delay of any scheme.

Index Terms—Ad hoc networks, delay tolerant networks, intermittent connectivity, routing.

I. INTRODUCTION

INTERMITTENTLY connected mobile networks (ICMN) are mobile wireless networks where most of the time there does not exist a complete path from a source to a destination, or such a path is highly unstable and may change or break while being discovered. There are many real networks that follow this model. Examples include wildlife tracking and habitat monitoring sensor networks [1], military networks [2], vehicular ad hoc networks (VANETs) [3], pocket switched networks (PON) [4], networks for low-cost Internet provision to remote communities [5], etc. In these networks, intermittent connectivity might arise due to sparseness [5], [6], nodes powering down to conserve energy [1], high mobility [3], or even for covertness [2]. Intermittently connected mobile networks

belong to the general category of Delay Tolerant Networks (DTN) [7], that is, networks where incurred delays can be very large and unpredictable.

Conventional Internet routing protocols (e.g., RIP, OSPF) as well as ad hoc network routing schemes, such as DSR, AODV, etc. [8], assume that a complete path exists between a source and a destination, and try to discover minimum cost paths before any useful data is sent. Since no such end-to-end paths exists most of the time in ICMNs, such protocols would fail in this context. However, this does not mean that packets can never be delivered under intermittent connectivity. Over time, different links come up and down due to node mobility (or other reasons). If the sequence of connectivity graphs over a time interval are overlapped, then an end-to-end path might exist. This implies that a message could be sent over an existing link, get buffered at the next hop until the next link in the path comes up, and so on and so forth, until it reaches its destination.

This imposes a new model for routing. Routing here consists of a sequence of independent, local forwarding decisions, based on current connectivity information and predictions of future connectivity information. Furthermore, node mobility often needs to be exploited in order to deliver a message to its destination, which is why this model is usually referred to as “mobility-assisted routing” (other names include “encounter-based forwarding” and “store-carry-and-forward”). The idea is reminiscent of the work in [9]. However, *mobility there is exploited in order to improve capacity, while here it is used to overcome the lack of end-to-end connectivity.*

Depending on the number of copies of a single message that may coexist in the network, one can define two major categories of mobility-assisted routing schemes. In *single-copy* schemes, there’s only one node in the network that carries a copy of the message at any given time. We call this node the “custodian” of the message. When the current custodian forwards the copy to an appropriate next hop, this becomes the message’s new custodian, and so on and so forth until the message reaches its destination. On the other hand, *multiple-copy* (or *multi-copy*) routing schemes may generate multiple copies of the same message, which can be routed independently for increased robustness.

The majority of routing schemes proposed for ICMNs are flooding-based, and, therefore, multi-copy in nature [1], [10], [11]. Despite their increased robustness and low delay, flooding-based protocols consume a high amount of energy, bandwidth, and memory space (all scarce resources for most low-cost wireless devices) [1], [10], [12]. Further, under high traffic loads they suffer from severe contention and message drops that can significantly degrade their performance and scalability [12], [13]. These shortcomings often make such

Manuscript received September 18, 2005; revised July 8, 2006; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Das. This material is based upon work supported by the National Science Foundation under Grant CNS-0520017.

T. Spyropoulos was with INRIA, Sophia Antipolis, France, and is now with ETH, Zurich, 8092 Zurich, Switzerland (e-mail: spyropoulos@tik.ee.ethz.ch).

K. Psounis and C. S. Raghavendra are with the Department of Electrical Engineering and Computer Science, University of Southern California, Los Angeles, CA 90089, USA (e-mail: kpsounis@usc.edu; raghu@usc.edu).

Digital Object Identifier 10.1109/TNET.2007.897962

algorithms inappropriate for energy-constrained and bandwidth-constrained applications, which is commonly the case in wireless networks. Consequently, it is desirable to design efficient single-copy routing schemes for many resource-constrained ICMNs. Additionally, single-copy schemes constitute the building blocks of multi-copy schemes. In many multi-copy schemes a number of copies are generated, each of which is routed independently using a single-copy algorithm [12]. For this reason, it is important to have a good understanding of the tradeoffs involved in single-copy routing, in order to design efficient multi-copy schemes, as well.

With this in mind, we perform in this paper a thorough investigation of single-copy routing for intermittently connected mobile networks. (In [14] we study the same problem using multi-copy approaches.) We present a number of increasingly “smart” schemes, exposing their individual advantages and shortcomings, and demonstrate that competitive performance can often be achieved without the overhead and logistics of using redundant copies. The champion algorithm of our study turns out to be one that combines the simplicity of a simple random policy, which is efficient in finding good leads towards the destination, with the sophistication of utility-based policies that efficiently follow good leads. Finally, we propose an analytical framework to evaluate the performance of any routing scheme in the context of ICMN. Using this framework we derive lower and upper bounds on the expected delay of any single-copy routing scheme (these are actually bounds for multi-copy schemes, as well). We also use our framework to analyze the expected delivery delay of all the single-copy algorithms presented.

In the next section we go over some existing related work. Section III describes a number of single-copy routing algorithms, including our proposed solution and an optimal scheme. Then, in Section IV we present our analytical framework, and use it to evaluate the expected performance of the routing schemes presented. Section V provides simulation results where the performance of all the strategies is compared. Finally, Section VI concludes the paper and gives some directions for future work.

II. RELATED WORK AND CONTRIBUTIONS

Although a large number of routing protocols for wireless ad hoc networks have been proposed [8], [15], traditional routing protocols are not appropriate for networks that are sparse and disconnected. The performance of such protocols would be poor even if the network was only “slightly” disconnected. To see this, note that the expected throughput of *reactive* protocols is connected with the average path duration PD and the time to repair a broken path t_{repair} with the following relationship: $throughput = \min\{0, DataRate(1 - (t_{repair}/PD))\}$ [16]. Node mobility leads to frequent disconnections, reducing the average path duration significantly. Consequently, in most cases t_{repair} (at least $2\times$ the optimal delay) is expected to be larger than the path duration, which implies that the expected throughput will be close to zero. *Proactive* protocols, on the other hand, would declare lack of a path, or result into a deluge of topology updates.

One approach to deal with very sparse networks or connectivity “disruptions” [2] is to reinforce connectivity on demand, by bringing for example additional communication resources into the network when necessary (e.g., satellites, UAVs, etc.). Similarly, one could force a number of specialized nodes (e.g., robots) to follow a given trajectory between disconnected parts of the network in order to bridge the gap [17], [18]. In yet other cases, connectivity might be predictable, even though it is intermittent (e.g., Inter-planetary networks, IPN [19]). Traditional routing algorithms could then be adapted to compute shortest delivery time paths by taking into account future connectivity [5], [20]. Nevertheless, such approaches are orthogonal to our work; our aim is to study what can be done when connectivity is neither enforced nor predictable, but rather *opportunistic* and subject to the statistics of the mobility model followed by nodes.

Despite a significant amount of work and consensus existing on the general DTN architecture [7], there has not been a similar focus and agreement on DTN routing algorithms, especially when it comes to networks with opportunistic connectivity. The simplest possible approach is to let the source or a moving relay node (*DataMule*) carry the message all the way to the destination (*Direct Transmission*) [6]. Although this scheme performs only one transmission, it is extremely slow [9]. A faster way to perform routing in ICMNs, called *Epidemic Routing*, is to flood the message throughout the network [11]. This scheme is guaranteed to find the shortest path when no contention exists for shared resources like wireless bandwidth and buffer space. Yet, it is extremely wasteful of such resources. What is worse, in realistic scenarios where bandwidth, memory space, or energy resources might be scarce, the performance of flooding degrades significantly [10], [12], [13].

A number of approaches have been taken to reduce the overhead and improve the performance of epidemic routing [1], [10], [13], [21]–[23]. In [21] the authors examine a number of different schemes to suppress redundant transmissions and clean up valuable buffer space after a message has been delivered with epidemic routing. In [13], [22] a message is forwarded to another node with some probability smaller than one (i.e., data is “gossiped” instead of flooded). Finally, in [1] a simple method to take advantage of the history of past encounters is implemented in order to make fewer and more “informed” forwarding decisions. The concept of history-based or utility-based routing is further elaborated in [10], [24]. Results from these studies indicate that using the age of last encounter with a node, when making a forwarding decision, results in superior performance than flooding. The concept of history-based routing has also been studied in the context of regular, connected wireless networks in [25]. Finally, it has also been proposed that *Network Coding* ideas could be useful to reduce the number of bytes transmitted [23]. Although all these schemes, if carefully tuned, can improve to an extent the performance of epidemic routing, they are still flooding-based in nature, and thus often exhibit the same shortcomings as flooding [14].

A different approach to significantly reduce the overhead of epidemic routing, while still maintaining good performance, is to distribute only a bounded number of copies [12], [21], [26], [27]. In a manner similar to the 2-hop scheme of [9], a copy is handed over to a fixed number of relays, each of which can then

deliver it only directly to the destination. Nevertheless, in many situations where node movement is strongly correlated or predominantly local, the performance of this scheme deteriorates [4], [14].

Despite the variety of existing approaches, the majority of them are multi-copy ones. Furthermore, the minority that deals with single-copy techniques only studies direct transmission [6] or some form of utility-based schemes in relatively different contexts [24], [25]. In this work, we perform a detailed inquiry into the problem space of single-copy routing, and show how to achieve competitive performance without using multiple copies. We look into how utility functions can be designed to fully take advantage of the “field” of past encounters, and propose a function that is shown to achieve up to an order of magnitude improvement in ICMNs over existing utility functions [25]. Finally, we propose a novel, hybrid routing scheme, which uses randomization when necessary to overcome some inherent shortcomings of utility-based forwarding.

In the theory arena, a large body of work has recently emerged trying to analyze the trade-offs involved between the *asymptotic* capacity and the *asymptotic* delay of the 2-hop scheme proposed in [9], and related schemes exploiting mobility [27]–[30]. Although asymptotic results provide valuable insight on the scalability of a given family of protocols, they often do not provide the necessary insight to design efficient and practical schemes. Furthermore, the majority of these works are concerned with delay in *connected* networks. In such networks, mobility is only used to reduce the number of transmissions. On the other hand, mobility in disconnected networks is an *intrinsic component* of the minimum delay. Furthermore, in connected networks the transmission range of each node has to scale with the number of nodes, in order to ensure connectivity, making all related analytical results strictly a function of the number of nodes [27], [28]. Here, we’re interested in a much wider range of connectivity scenarios, where transmission range, number of nodes, and network size are *independent* parameters, whose individual effect on performance our analytical framework aims at quantifying.

Also, in the context of disconnected networks, most existing analytical efforts concern the performance of epidemic routing or other multi-copy schemes [14], [21], [22], [26]. To the best of our knowledge, the only prior analytical work for single-copy schemes is on direct transmission [6] and some asymptotic results regarding utility-based schemes [25], [31]. Finally, in many existing studies, some parameters of the proposed model (e.g., node inter-meeting times) need to be acquired from simulation traces for each and every scenario [21], [22]. Here, we expand our framework from [32] to evaluate the delivery delay of all the single-copy algorithms examined, as well as to derive lower and upper bounds on the expected delay achievable by any scheme in ICMNs.

III. SINGLE-COPY ROUTING STRATEGIES

In this section, we explore the problem space of single-copy routing in ICMNs. Our problem setup consists of a number of nodes moving independently according to some stochastic mobility model. Additionally, we assume that the network is disconnected at most times, and that transmissions are faster

than the node movement (a reasonable assumption with modern wireless devices¹). Also, each node can maintain a timer for every other node in the network, which records the time elapsed since the two nodes last encountered each other (i.e., came within transmission range). These timers are similar to the *age of last encounter* in [25], and are useful, because they contain indirect (relative) location information. However, note that not every routing scheme requires these timers. Also, we assume that this is the *only* information available to a node regarding the network (i.e., no explicit location info, etc.). Finally, we assume that nodes emit a beacon signal, possibly periodically, that advertises their presence. In practice, when another node senses this beacon, the two nodes establish a relationship (as for example in Bluetooth pairing [33]) by exchanging IDs and other relevant information like timer values. We refer to this as an “encounter”.²

We will now look into a number of increasingly “smart” routing strategies. We believe that these are fairly representative of different approaches one might take for the problem in hand. Each routing algorithm decides under what circumstances a node, currently holding the single message copy, will hand it over to another node it encounters. The goal is that each forwarding step should, on the average, result in progress of the message towards its destination. (Due to space limitations, all the proofs for this section can be found in [34].)

Direct Transmission: The simplest possible routing scheme imaginable is the following: *a node A forwards a message to another node B it encounters, only if B is the message’s destination.* This scheme is trivial, but it has the advantage of performing only a single transmission per message. It has been considered in some previous work [6], [9], and its expected delivery delay is an upper bound on the expected delay of any routing scheme. It will therefore serve as our baseline.

Randomized Routing Algorithm: The first nontrivial routing algorithm that we will look at is a randomized forwarding algorithm, where the current message custodian hands over the message to another node it encounters with probability $p \in (0, 1]$. Further, in order to avoid a message constantly jumping back and forth between two nodes within range, we assume that, when a node receives a message, it is not allowed to send the message back to the node it received it from, for a given amount of time (the two nodes are tagged as “coupled” [35] until a timer expires).

Lemma 3.1: When all nodes move according to a stochastic mobility model whose expected meeting time is a concave function of distance, forwarding a message results in a reduction of the expected delivery time of the message to its destination.

Lemma 3.1 states that even this simple routing strategy results in expected progress at each forwarding step (i.e., locally), for a number of mobility models (e.g., Random Walk, Random Waypoint [8]). This result might be slightly counterintuitive, but can

¹Assume, for example, that a node has a range of 100 m and a radio of 1 Mbps rate. Then, it could send a packet of 1 KB at a distance of 100 m in only 8 ms. Even if that node is a fast moving car with a speed of say 65 mph, it could carry the same packet at a mere distance of less than 1 m in the same 8 ms.

²Although the frequency with which the beacon is emitted plays a role on when nodes encounter each other, here we will assume for simplicity that nodes instantaneously “see” each other when they come within range.

be explained by the fact that transmissions are faster than node movement. Nevertheless, this progress is marginal, especially when far from the destination.

Utility-Based Routing With 1-Hop Diffusion: Randomized Routing does not take advantage of the only information available to each node regarding the network, that is, the last encounter timers. Position information regarding different nodes gets indirectly logged in the last encounter timers, and gets diffused through the mobility process of other nodes. For many noncontrived (or nonadversarial) mobility models, it can be shown that a smaller timer value on average implies a smaller distance from the node in question. Therefore, we can define a utility function based on these timers, which indicates how “useful” a node might be in delivering a message to another node. A gradient-based scheme can then be used to deliver a message to its destination, as has been noted in [10], [25]. This scheme will try to maximize the utility function for this destination.

Definition 3.1 (Utility-Based Routing): Let $\tau_i(j)$ denote the time elapsed since node i last saw node j . Let further each node i maintain a utility function $U_i(\cdot)$ for all nodes, where $U_i(\cdot)$ is a monotonically decreasing function of the respective last encounter timer $\tau_i(\cdot)$, and $U_i(i) \geq U_i(j)$, $\forall i, j$. Then, a node A forwards to another node B a message destined to a node D , if and only if $U_B(D) > U_A(D) + U_{th}$, where U_{th} (utility threshold) is a parameter of the algorithm.

Lemma 3.2: Let all nodes move according to a stochastic mobility model for which $\tau_A(D) > \tau_B(D) \Leftrightarrow E[d_{AD}] > E[d_{BD}]$, (d_{ij} denotes the distance between two nodes). Let further a node A carrying a message for a node D encounter another node B , for which $U_A(D) < U_B(D)$. Then, forwarding the message to B results in reduction of the expected delivery time of the message to its destination, which is *at least* as large as that achieved by Randomized Routing.

The above lemma, as in the case of Randomized Routing, holds for a number of mobility models like Random Walk, Random Waypoint, Community-based mobility [36], etc.

Utility-Based Routing With Transitivity: Despite making better forwarding decisions than randomized ones, the previous scheme suffers from a “slow start” phase, which is more manifested in large networks. In a large network, where source and destination are usually far, almost all nodes around the source will have moved long enough to get “decoupled” from the destination [35]. Thus, they will not have a high enough utility to become next hops. Additionally, if it happens that the few nodes around the message custodian last met the destination before the custodian did, the custodian will probably have to wait a long time until it moves close to the destination again (even if a connected path to the destination existed). The reason for this inefficiency is that each node updates its utility function for the destination only when it encounters that destination. Location information takes a very long time this way to get diffused throughout the network, and by the time such information does get diffused it has become obsolete.

To deal with this problem we propose the use of “transitivity” when updating the utility function. *When node A sees node B often, and node B sees node C often, A may be a good candidate to deliver a message to C (through B), even if A rarely sees*

C. Therefore, when A encounters node B , it should also update (increase) its utility for all nodes for which B has a high utility.³ Although the idea of using transitivity is not entirely new (e.g., see [10]), the transitivity function needs to be carefully chosen, in order to actually improve performance.⁴ From an information theoretic perspective, this transitivity effect should successfully capture the amount of uncertainty resolved regarding the position of the destination, when a node is encountered that has some additional (i.e., more recent) information for that destination. For reasons that will become clear in Section IV, we propose the use of the following transitivity function that we apply directly to the last encounter timers:

Definition 3.2 (Timer Transitivity): Let a node A encounter a node B at distance d_{AB} . Let further $\bar{t}_{mm}(d)$ denote the expected time it takes a node to move a distance d under a given mobility model. Then: $\forall j \neq B : \tau_B(j) < \tau_A(j) - \bar{t}_{mm}(d_{AB})$, set $\tau_A(j) = \tau_B(j) + \bar{t}_{mm}(d_{AB})$.

Seek and Focus Routing—A Hybrid Approach: Although the use of transitivity does alleviate the slow start phase, if nodes move fast enough, even transitivity might not be able to diffuse utility information promptly throughout the network. Additionally, utility values of nodes can be seen as a time-varying utility field with the global maximum at the destination. Since a greedy forwarding approach is used, the message often gets stuck at local maxima for some time.

In order to deal with these shortcomings, we propose a hybrid routing protocol, called “Seek and Focus”, which aims to combine the best of both worlds. It can escape the slow-start phase and local maxima of utility-based protocols, while still taking advantage of the higher efficiency of utility-based forwarding. Initially it looks around greedily for a *good lead* towards the destination using randomized routing, and then uses a utility-based approach to follow that lead efficiently. Additionally, it uses a procedure reminiscent of the “peripheral routing” in some position-based routing protocols [37] to prevent a message from getting stuck for a long time at local maxima of utility.

Definition 3.3 (Seek and Focus (Hybrid)): Seek and Focus consists of the following alternating phases:

- *seek phase:* If the utility of the current custodian is below a predefined threshold U_f (“focus threshold”), perform randomized forwarding with parameter p to quickly search nearby nodes;
- *focus phase:* if a node with a utility above U_f just received the message, then (i) this node resets a timer t_{focus} to T_{local} and starts counting down, and (ii) it performs utility-based forwarding (i.e., looks for a node with a higher utility than its own);
- *re-seek phase:* if t_{focus} expires then (i) set timer t_{seek} to T_{per} and U_f to the current utility value, and (ii) perform randomized routing until either a node with utility higher

³In practical situations, each node would only maintain a cache of the most recent nodes that it has encountered, in order to reduce the overhead of book-keeping and the amount of timer data exchanged when two nodes encounter each other. We expect this to have only a small, if any, effect on performance, as large timer values contain little information anyway.

⁴Simulations we performed have shown us that heuristic transitivity functions like the one proposed in [10] only increase the number of forwarding steps per message, but decrease the “quality” of each step; the protocol’s behavior starts resembling that of randomized routing.

than U_f is found or until t_{seek} expires; then, reset U_f to its default value and return to seek ($utility < U_f$) or focus phase ($utility > U_f$).

An “Oracle-Based” Optimal Algorithm: The algorithm that minimizes the expected delivery delay is aware of all future movement, and, thus, it is an “oracle-based” algorithm. Based on this knowledge, it computes the *optimal set of forwarding decisions* (i.e., time and next hop), which delivers a message to its destination in the minimum amount of time.

The “oracle-based” algorithm cannot be implemented in reality, when connectivity is opportunistic. It provides an *offline solution to an inherently online problem*, and, thus, its delay will serve as a lower bound on the delay of any routing strategy. (Note that, when future connectivity is known, it could be possible to implement this scheme, albeit with considerable overhead [5].) Finally, notice that, under the assumption of infinite buffer space and bandwidth, flooding (i.e., epidemic routing) achieves this minimum delay.

IV. PERFORMANCE ANALYSIS

In this section, we will analyze the expected delay of all routing schemes presented, under intermittent connectivity. Since message transmissions in this context occur only when nodes meet each other, *the time elapsed between such meetings is the basic delay component*. Therefore, in order to be able to evaluate the performance of any mobility-assisted scheme, it is necessary to know the statistics of encounter times between nodes, called *hitting* or *meeting* times. These are the times until a node, which, say, just received a message, *first* encounters a given other node that can act as a relay.

Definition 4.1 (Hitting and Meeting Time): Let a node i move according to mobility process “MM”, and starting from its stationary distribution at time 0. Let further its position at time t be $X_i(t)$. Then,

- i) If j is a static node with uniformly chosen position X_j , then the *expected hitting time* under mobility model MM is $ET_{mm} = \min_t \{t : \|X_i(t) - X_j\| \leq K\}$.
- ii) If j is a mobile node also starting from its stationary distribution, then the *expected meeting time* between the two nodes is $EM_{mm} = \min_t \{t : \|X_i(t) - X_j(t)\| \leq K\}$.

These constitute the basic component in the delay expression of any scheme, and they largely vary depending on the specific mobility model in hand. In [36] we have calculated the meeting time for “epoch-based” mobility models, like Random Waypoint, Random Direction, and a more realistic mobility model called Community-based mobility. Here, we will start by complementing this analysis with the calculation of hitting and meetings times for independent Random Walks. We will then analyze the delay of the single-copy strategies schemes presented. Due to lack of space, we will focus in most cases on Random Walk mobility. However, the general methodology we present applies to other mobility models as well, specifically models that have approximately exponentially distributed meeting times. These include many popular mobility models like Random Direction, and Random Waypoint, as well as synthetic models based on these, like Community-based mobility [21], [36]. Throughout the section we will make the following assumptions:

TABLE I
NOTATION

N	size of network area
M	number of nodes
K	transmission range
$X_i(t)$	The position of node i at time t , $(x_i(t), y_i(t))$; $X_i = (x_i, y_i)$ if the node is static.
ET_{mm}	expected hitting time under “MM” (randomly chosen positions)
EM_{mm}	expected meeting time under “MM” (randomly chosen positions)
$E_i T_j$	expected time until a node starting at position i first hits position j
$E_i T_j(K)$	expected time until a node starting at position i first comes within range K of position j
$E_i T_A$	expected time until a node starting at position i first hits a subset of positions A
$E_{\pi(s)} T_j$	expected time until a node starting from distribution $\pi(s)$ hits position j
$E_{\pi(s)} T_j(K)$	expected time until a node starting from distribution $\pi(s)$ comes within range K of position j
EM_{ij}	expected time until two nodes, starting at positions i and j , respectively, first meet each other
T_j^+	time until a node starting from position j first returns back to position j
ED_{rp}	expected delay of routing protocol “RP”

Network: M nodes move on a $\sqrt{N} \times \sqrt{N}$ two-dimensional torus (we choose torus for symmetry; we have run simulations that show that performance in bounded networks is similar). Each node can transmit up to distance $K \geq 0$ meters away, where K/\sqrt{N} is smaller than the value required for connectivity [38], and each message transmission takes one time unit.

Last Encounter Timers: Every node i maintains a timer $\tau_i(j)$ for every other node j it has encountered. Initially all $\tau_i(j)$ are set to ∞ . $\tau_i(i) = 0, \forall i$ at every time instant. When node i encounters node j it sets $\tau_i(j)$ to 0, and increases this timer for every time unit elapsed.

Contention: Throughout our analysis we assume that bandwidth and buffer space are infinite. In other words, we assume that there is no contention for these resources. Although contention is an important factor for multi-copy schemes, we argue that it is significantly less of an issue for single-copy schemes that perform much fewer transmissions.⁵

Despite some of these assumptions being somewhat unrealistic (e.g., IID node mobility), we will show in the simulations section that our analysis can still capture the performance behavior of different protocols, even in more realistic scenarios. Finally, in most theorems and proofs, we use the same notation as in [35]. However, we summarize most of our notation in Table I.

A. Direct Transmission—An Upper Bound on Delay

We will first calculate the delay of Direct Transmission ED_{dt} , assuming that nodes perform independent random walks on a 2-D grid (i.e., discrete model). This delay equals the expected meeting time EM_{RW} between two nodes, and

⁵It is important to note that, by Little’s law, if the input traffic increases enough, single-copy schemes would also eventually face problems with contention. However, they “break” much later than multi-copy schemes do, while for the traffic loads we’re interested at they essentially operate contention-free. A first effort to model contention for the wireless media and buffer space can be found in [22], [39].

it constitutes a lower bound on the expected delay of any other mobility-assisted routing strategy, under Random Walk mobility. Furthermore, we will use here Manhattan distance $d_{ij} = |x_i - x_j| + |y_i - y_j|$ to measure proximity between two nodes (or their positions) i and j , since this corresponds better, analytically, to the discrete nature of random walks on grids. If Euclidean distance was used instead, respective delays would only be slightly smaller, due to the slightly larger coverage for the same range K .

When transmission range $K = 0$ (i.e., nodes encounter each other only when they occupy the same position), the hitting and meeting times for random walks can be found in [35]. In the general case where $K \geq 0$, a node i encounters another node j when i first hits the subset of grid positions $A(K) = \{a : |x_a - x_j| + |y_a - y_j| = K\}$ (this is the ‘‘border’’ of the area within range of j). Therefore, the expected delay of direct transmission $ED_{dt} = EM_{RW}$ depends on the expected hitting time $E_{\pi}T_{A(K)}$, where π denotes the stationary distribution on the grid.

Before we proceed with the calculation of $E_{\pi}T_{A(K)}$, we prove a result that we will use there, regarding hitting times on subsets. Its proof is based on the electric network analogy for random walks on graphs. This is a known method that somewhat simplifies the combinatorics involved when calculating the numbers of paths that a random walk can follow, in a given problem setting [35], [40].

Lemma 4.1: Let j be a position on the grid, and let $A(K)$ denote the subset of all positions a , such that $|x_a - x_j| + |y_a - y_j| = K$. Let further $T_{A(K)}$ denote the time until a random walk starting from j first hits $A(K)$. The probability $P(T_{A(K)} < T_j^+)$ that a walk starting from j hits $A(K)$ before it returns to j is then given by

$$P(T_{A(K)} < T_j^+) = \frac{2^K - 1}{2^{K+1} - K - 2}. \quad (1)$$

Proof: See the Appendix. ■

Theorem 4.1 first calculates the expected hitting time on the border of the coverage area of j , $E_{\pi}T_{A(K)}$, as: $E_{\pi}T_j$, the expected hitting time to grid position j (larger than $E_{\pi}T_{A(K)}$), minus $E_{\pi(A)}T_j$, the expected time to go from $A(K)$ to j . Then, based on this, it calculates $ED_{dt} = EM_{RW}$.

Theorem 4.1: Let us pick a position j uniformly on the torus, and let ET_{RW} denote the expected *hitting time* until a random walk, starting from the stationary distribution, comes within range $K \geq 0$ of j . Further, let EM_{RW} denote the expected *meeting time* until two random walks, starting from the stationary distribution, come within range K . Then, for large N ,

$$i) \quad ET_{RW} = N \left(0.34 \log N - \frac{2^{K+1} - K - 2}{2^K - 1} \right) \quad (2)$$

$$ii) \quad ED_{dt} = EM_{RW} = \frac{1}{2} ET_{RW}. \quad (3)$$

iii) The distributions of T_{RW} , M_{RW} have exponential tails.

Proof:

i) Let us assume that the random walk starts outside of j 's range. Then, $ET_{RW} = E_{\pi}T_{A(K)}$. Now let π_A denote the stationary distribution on set $A(K)$, that is,

$$\forall i \in A(K), \pi_A(i) = \pi(i)/\pi(A),$$

where $\pi(A) = \sum_{i \in A(K)} \pi_i$. Then, it holds that:

$$E_{\pi}T_j = E_{\pi}T_{A(K)} + E_{\pi_A}T_j.$$

It is known from [35] that, for $N \rightarrow \infty$,

$$E_{\pi}T_j = cN \log N,$$

where $c = 0.34$. This is actually ET_{RW} for $K = 0$. Also, according to [6], [41] this results is quite accurate for $N > 25$.

We're now only missing $E_{\pi_A}T_j$ to acquire $ET_{RW} = E_{\pi}T_{A(K)}$. Let P_A denote $P(T_{A(K)} < T_j^+)$ see Lemma 4.1), and let us further express the *expected* first return time of a walk starting at j , $E_jT_j^+$, as a weighted average of two cases:

- (a) the walk reaches $A(K)$ before it returns back to j (probability P_A);
- (b) the walk returns to j without reaching $A(K)$ (probability $1 - P_A$).

Thus,

$$E_jT_j^+ = P_A \cdot (E_jT_{A(K)} + E_{\pi_A}T_j) + (1 - P_A)E_j [T_j^+ | T_{A(K)} > T_j^+],$$

where $E_j [T_j^+ | T_{A(K)} > T_j^+]$ is the expected first return time, *conditioned on the fact that $A(K)$ was not reached*. Using Kac's formula [35] we get that

$$E_jT_j^+ = 1/\pi(j) = N.$$

Furthermore, according to [35] it holds that

$$E_jT_{A(K)} \leq |A(K)|^2 = O(K^2),$$

$$E_j [T_j^+ | P(T_{A(K)} > T_j^+)] = O(K).$$

Hence, the previous expression can be rewritten as

$$P_A \cdot E_{\pi_A}T_j = N - f(K) - g(K),$$

with

$$f(K) = P_A \cdot E_jT_{A(K)},$$

$$g(K) = (1 - P_A)E_j [T_j^+ | T_{A(K)} > T_j^+],$$

being two functions depending only on K (note that P_A is also a function of K only). Consequently, when $N \gg K$ (as expected in ICMNs), the above equation implies that

$$P_A \cdot E_{\pi_A}T_j \approx N.$$

Replacing P_A from (1) we get that

$$E_{\pi_A}T_j = \left(\frac{2^{K+1} - K - 2}{2^K - 1} \right) N.$$

Finally, replacing $E_{\pi_A}T_j$ and $E_{\pi}T_j$ in our initial expression for $E_{\pi}T_j$ we get $E_{\pi}T_{A(K)} = ET_{RW}$.

- ii) The proof is a straightforward extension of the proof for the case of $K = 0$ [35].
- iii) It is proven in [35] that, for random walks, hitting times on subsets have an exponential tail. ■

Although point (iii) of the above Theorem only proves an exponential tail, for most values of interest the distribution of the meeting time is closely approximated by an exponential with average EM_{RW} . We will provide further evidence for this in Section IV-E.

B. “Oracle-Based” Optimal Algorithm—A Lower Bound on Delay

The following coloring problem analog is used to analyze the expected performance of the optimal algorithm: The randomly chosen source node is colored *red* and all other nodes are colored *blue*; whenever a red node encounters a blue node, the latter is colored red too. It is evident that the expected time until the destination node is colored red is equal to the expected message delivery time of the optimal algorithm. It is also evident by this coloring analog why flooding will find the same path as the oracle scheme, when buffer capacity and bandwidth are unlimited.

The following theorem calculates the expected time until the destination node is colored red. It uses the expected meeting time, under the given mobility model, to calculate the expected time until a new node is colored red, and then estimates the total number of coloring steps necessary.

Theorem 4.2: Let us assume that all nodes move according to a stochastic mobility model whose meeting time probability distribution has an exponential tail, and an average value of EM_{mm} . Then, the expected message delivery time of the optimal algorithm $ED_{opt}^{(mm)}$ is given by

$$ED_{opt}^{(mm)} = \frac{H_{M-1}}{(M-1)} EM_{mm}, \quad (4)$$

where H_k is the k^{th} Harmonic Number, i.e., $H_k = \sum_{i=1}^k (1/i) = \Theta(\log k)$.

Proof: We assume all M nodes start from the stationary distribution and remain so at all times.⁶ Let us further assume that, at some time instant, we have n red nodes and $M-n$ blue nodes, and let us pick a red node i and a blue node j . Further, let M_{ij} denote the meeting time between two nodes i and j , under the mobility model in hand. If M_i denotes the meeting time of i with any of the $M-n$ blue nodes, then $M_i = \min_j(M_{ij})$. Finally, let $M^{(n)}$ denote the time until any of the red nodes meets any of the blue ones, when there are n total red nodes. Then, $M^{(n)} = \min_i(M_i) = \min_{ij}(M_{ij})$. If all M_{ij} are IID exponential random variables with average EM_{mm} , then $M^{(n)}$ is also an exponential random variable with average $EM_{mm}/n(M-n)$. Finally, since we have started with 1 red node, the time until all nodes are colored red is given by $\sum_{n=1}^{M-1} M^{(n)}$, whose expected value can be calculated by

$$\sum_{m=1}^{M-1} \frac{EM_{mm}}{n(M-n)}.$$

This is the expected time until *all* nodes are colored red. However, the destination may be colored red in any of the $M-1$ total coloring steps with equal probability. Consequently, the

⁶Stationarity is a desirable property for mobility models used in simulations, and can be achieved for models that can be represented with an appropriate Markov Chain, if the nodes are started from stationarity [42].

expected coloring time for the destination is the average of all these cases:

$$\frac{1}{M-1} \sum_{l=1}^{M-1} \sum_{n=1}^l \frac{EM_{mm}}{n(M-n)} = EM_{mm} \frac{H_{M-1}}{M-1}. \quad \blacksquare$$

It is important to note that (4) is generic. By plugging into the equations the appropriate meeting time value EM_{mm} (e.g., EM_{RW} from Theorem 4.1), we can calculate the optimal delay for the respective mobility model.

C. Performance of Randomized Routing

In the case of the randomized routing algorithm, the single message copy performs a random walk on the dynamically changing connectivity graph.

Definition 4.2 (Message’s Random Walk): Let a message be routed according to the randomized routing algorithm. Let also p_e be the probability that there is at least one more node within the range of the current custodian j , and recall that p is the probability of forwarding. Then the message’s movement can be modelled as follows: at every step, (i) with probability $1-p \cdot p_e$ it moves according to the respective mobility model (e.g., in the case of pure random walks, it moves to neighboring positions with probability 1/4); (ii) with probability $p \cdot p_e$ it jumps to a position in set $N_j(K) = \{i : d_{ij} \leq K\}$, which is the subset of all positions that are within range of j .

The following Lemma calculates the probability of encounter p_e and the average transmission distance $\bar{d}_{tx}(K)$ that a message covers if it is transmitted, when we assume Manhattan distance.

Lemma 4.2: Let a message perform a random walk according to Definition 4.2, in a network consisting of M nodes performing independent random walks on an $\sqrt{N} \times \sqrt{N}$ torus. Then, the probability of encounter p_e of the message at any time is given by

$$p_e = 1 - \left(1 - \frac{2K^2 + 2K + 1}{N}\right)^{M-2}. \quad (5)$$

Additionally, the average transmission distance is given by

$$\bar{d}_{tx}(K) = \frac{K(8/3 + 2K + 4/3K^2)}{2K^2 + 2K + 1}. \quad (6)$$

Proof: Both equations can be derived using elementary probability theory and combinatorics. \blacksquare

We are now ready to analyze the performance of the randomized algorithm. Theorem 4.3 provides an upper bound on its delivery delay.⁷

Theorem 4.3: Let all nodes perform independent random walks, and let each node have a transmission range $K \geq 1$. Then, the expected message delivery time of the randomized algorithm $ED_{rnd}^{(RW)}$ is given by

$$ED_{rnd}^{(RW)} \leq \frac{ET_{RW}}{2 - p \cdot p_e + p \cdot p_e \bar{d}_{tx}(K)}. \quad (7)$$

⁷It is important to note that (5), and thus Theorem 4.3, ignore the effect of “coupled” nodes. Although it is not too difficult to take coupling into account, its effect becomes negligible as M increases, and we choose to ignore it for our analysis.

Proof: Let two independent random walks, S and D , start from positions x and y , respectively. Walk D models the destination's movement, which is a pure random walk on the torus, while walk S models the message movement and moves according to Definition 4.2. When no transmission occurs, S (the message) will move only to a neighboring position, while it makes a jump of average length $\bar{d}_{tx}(K)$ when a transmission occurs. Consequently, the average length of a jump of S is given by

$$\bar{d}(K) = (1 - p \cdot p_e) \cdot 1 + p \cdot p_e \bar{d}_{tx}(K).$$

Now, consider another random walk S' , starting also at x , which whenever walk S jumps d positions far, it performs d independent single step movements (i.e., performs d pure random walk steps), *within one time unit*. It is evident that S will cover (i.e., bring within its transmission range) a higher number of new positions than S' , at every step, and therefore S is expected to meet the destination node D faster than S' .⁸ It is more convenient to analyze the behavior of S' .

Let $(X(t), Y(t))$ be the positions of (S', D) after t moves, and let us define the function

$$f(x, y) = E_x T_y(K) - E_\pi T_y(K).$$

Finally, consider

$$W_t \equiv (\bar{d}(K) + 1)t + f(X(t), Y(t))$$

and define as M_{xy} the first meeting time between S' and D (i.e., assuming that S' started from position x and D from position y).⁹ It is easy to verify that

$$(W_t; 0 \leq t \leq M_{xy})$$

is a martingale [43] (using similar arguments as in [35, Ch. 3, Proposition 3]). According to the optional stopping theorem [43] $EW_0 = EW_{M_{xy}}$. This means that

$$E_x T_y(K) - E_\pi T_y(K) = (\bar{d}(K) + 1) EM_{xy} + Ef(X(M_{xy}), Y(M_{xy})).$$

However, by definition $X(M_{xy}) = Y(M_{xy})$ and therefore

$$E_{X(M_{xy})} T_{Y(M_{xy})}(K) = 0.$$

Consequently,

$$E_x T_y(K) - E_\pi T_y(K) = (\bar{d}(K) + 1) EM_{xy} - E_\pi T_{Y(M_{xy})}(K).$$

Further,

$$E_\pi T_y(K) = ET_{RW}, \quad \forall y$$

⁸Assume $d = 2$, for example. At the end of a step, walk S is distance d far from its previous position, and has covered approximately $d \cdot (2K + 1)$ new positions. On the other hand, after S' has performed $d = 2$ independent random walk steps at the same time, it might have ended up at its previous position with probability $1/4$ (covering only $2K + 1$ new area in the intermediate step).

⁹ S' performs $\bar{d}(K)$ steps per time unit, and D performs an additional step at the same time unit, giving a total of $\bar{d}(K) + 1$ steps.

and given by (2). Consequently,

$$E_x T_y = (\bar{d}(K) + 1) EM_{xy}.$$

Finally, for uniformly chosen x, y , $E_x T_y = ET_{RW}$, giving us

$$EM_{xy} = ET_{RW} / (\bar{d}(K) + 1).$$

Also, for uniformly chosen x, y , EM_{xy} is the expected meeting time of S' and D , which is an upper bound on the delivery time of the randomized algorithm (i.e., meeting time of S and D). ■

A similar methodology has been used to calculate the expected delay of Randomized Routing under the Random Direction mobility model [36], and could also be used for other mobility models with exponential meeting times.

D. Performance of Schemes Using Utility Functions

In utility-based forwarding the message will go with the node that is expected to be "closer" to the destination. Thus, unlike Randomized Routing, in the case of Utility-based Routing the message movement has a bias towards the destination at every step. Furthermore, the amount of this bias depends on how far the message is from the destination. The system has now memory (state) that forces us to depart from the previous model, where each step was statistically equivalent. Therefore, in order to calculate the expected delay of schemes that use utility functions we will model the message's movement as a Markov Chain whose state is the message's distance r from the destination.

We use this model to analyze the expected delay of Utility-based routing and Seek and Focus routing for the case where no transitivity is used in the utility function, and nodes perform independent random walks. Transitivity complicates matters even further since it introduces multi-dimensional memory, where the random variable is not just the past trajectory of a node compared to that of the destination, but rather the combined trajectories of all nodes together. For this reason, we will only provide here some theoretical insight on the effect of transitivity on performance and illuminate our choice for the utility function of Definition 3.2.

Utility-Based Routing Without Transitivity: Without loss of generality, we assume $U_X(Y) \equiv \tau_X(Y), \forall X, Y$, and thus the routing scheme aims at *minimizing* $U_X(\text{destination})$. We can calculate the transition probabilities for the corresponding Markov chain as follows: Let the current message custodian lie at distance r from the destination (i.e., state r of the chain). We know that at any time, a node is within range with probability p_e (see Section IV-C). If that node has a higher utility value than the message custodian, then the message gets forwarded. With some probability, let $p_{tx}(r \rightarrow r - i)$, this message ends up at state $r - i$ (i.e., closer to the destination), and with some probability, let $p_{tx}(r \rightarrow r + i)$, it ends up at state $r + i$ (i.e., further away from the destination), where $i \leq K$. Finally, if no node is within range, or the node in range has a lower utility value than the current, the message performs a pure random walk step (carried by the node).

The following two Lemmas calculate $p_{tx}(r \rightarrow r \pm i)$. Lemma 4.3 first derives some useful probability density functions (pdf)

for a node performing a random walk on a torus. Specifically, if a node was chosen randomly (e.g., encountered at some point), (i) what might its distance from the destination be, (ii) what timer values might it have, and (iii) how would our knowledge about its distance change if we knew its timer value. These quantities are necessary for Lemma 4.4 to calculate next all $p_{tx}(r \rightarrow r \pm i)$.

Lemma 4.3: Let $u(r)$ denote the probability of a randomly chosen node be found at distance r from the destination, and $u(t)$ the probability that it has a timer value of t for the same destination. Further, let U denote the transition probability matrix for a graph that is the same as the $\sqrt{N} \times \sqrt{N}$ torus except that no transition from any state at distance $K+1$ (from the destination) to a state at distance K (or smaller) is allowed. Then:

$$i) \ u(r) = \begin{cases} 1/N, & r = 0 \text{ or } r = \sqrt{N}, \\ 4r/N, & 0 < r < \sqrt{N}/2, \\ 4r - 2/N, & r = \sqrt{N}/2, \\ 4(N-r)/N, & \sqrt{N}/2 < r < \sqrt{N}. \end{cases}$$

ii) $u(t)$ is exponential with average ET_{RW} ,

iii) $u(r|t) = U^t \cdot u_0$, where u_0 is a probability vector with value $1/(4(K+1))$ for all torus positions at distance $K+1$ and 0 for all other positions.

Proof: The proof can be found in the Appendix. ■

Lemma 4.4: Let us assume that a message is currently with a node A at distance r_A from its destination, and that another node B is encountered. Then, the probability $p_{tx}(r_A \rightarrow r_B)$ that B has a higher utility for the destination (i.e., message gets forwarded) and is at distance r_B is equal to

$$p_{tx}(r_A \rightarrow r_B) = \sum_{t_A} \sum_{t_B < t_A} \frac{u(r_A|t_A)u(t_A)u(r_B|t_B)u(t_B)}{u(r_A)u(|r_B - r_A| \leq K|t_B)}.$$

Proof: The proof can be found in the Appendix. ■

Theorem 4.4 uses the probabilities of Lemma 4.4 to define a Markov Chain that captures the message movement, and then uses this chain to calculate the expected delay of Utility-based routing.

Theorem: Let us define a Markov Chain with state r , and the following transition probabilities $p(r; r+i)$:

$$p(r; r+i) = \begin{cases} p_e p_{tx}(r \rightarrow r+i), & \text{if } |i| \leq K, |i| \neq 1, \\ \frac{1}{2} p_{rw} + p_e p_{tx}(r \rightarrow r+i), & \text{if } |i| = 1, \\ 0, & \text{otherwise.} \end{cases}$$

where $p_{tx}(r \rightarrow r \pm i)$ are given by Lemma 4.4, and $p_{rw} = [1 - p_e \sum_{|i| \leq K} p_{tx}(r \rightarrow r+i)]$.

Let further $\pi_u(r)$ denote the stationary distribution for this chain. Then, the expected delivery delay of Utility-based Routing $ED_{util}^{(RW)}$ is equal to $E_{\pi_u(r)} T_K$, the expected hitting time on state K (or smaller) of the chain, starting from $\pi_u(r)$.

Proof: According to the previous discussion, when a node is encountered at distance i from the current custodian the messages gets forwarded with probability $p_{tx}(r \rightarrow r+i)$. Hence, a transition should occur on the Markov chain from state r to state $r+i$ with probability $p(r; r+i) = p_e p_{tx}(r \rightarrow r+i)$, $\forall r, i$ such that $(-K \leq i \leq K)$ and $(0 \leq r+i \leq \sqrt{N})$; $p(r; r+i) = 0$ for all other r, i .

Further, if no node is within range, or the node encountered has a lower utility, then the message will not be transmitted.

It is easy to see that the former occurs with probability $1 - p_e$, while the latter with probability $p_e(1 - \sum_{|i| \leq K} p_{tx}(r \rightarrow r+i))$. Putting it altogether, with probability $p_{rw} = [1 - p_e \sum_{|i| \leq K} p_{tx}(r \rightarrow r+i)]$ the message only performs a pure random walk step. In that case, it moves with equal probability $1/2$ to state $r-1$ or to state $r+1$. Thus, a factor of $(1/2)p_{rw}$ needs to be added to all transition probabilities $p(r; r \pm 1)$.

Starting therefore from the stationary distribution $\pi_u(r)$, the expected delay until the message is delivered is equal to the time until the message hits a state smaller or equal to K on the chain. ■

The expected hitting time on the above one-dimensional chain can easily be calculated using the t -step transition matrix for the chain [43], giving us the expected delay of Utility-based Routing.

Seek and Focus Routing Without Transitivity: It is not too difficult to modify the above procedure to calculate the expected delay of Seek and Focus as well. We need to take into account here that when the current custodian has a low utility value, it does not look into the respective utility of the node encountered, but performs randomized routing instead. Thus, in the case of Seek and Focus $p_{tx}(r \rightarrow r+i)$ is different from the Utility-based forwarding case, and consists now of two terms: one that corresponds to the focus phase, and a second one that corresponds to the seek phase. (Note that, in order not to depart from the above Markov chain model, we will ignore in this analysis the effect of the “re-seek” phase of Seek and Focus.)

Lemma: Let us assume that a message is currently with a node A at distance r_A from its destination, and that another node B is encountered. Let us further assume that Seek and Focus Routing is implemented by both nodes, with a one-to-one utility function $U(\cdot)$, focus threshold U_f , and $t_f = U^{-1}(U_f)$. Then, the probability $p_{tx}(r_A \rightarrow r_B)$ that the message proceeds to a state r_B (i.e., is forwarded to B and that B is at distance r_B) is equal to:

$$p_{tx}(r_A \rightarrow r_B) = \sum_{t_A < t_f} \sum_{t_B < t_A} \frac{u(r_A|t_A)u(t_A)u(r_B|t_B)u(t_B)}{u(r_A)u(|r_B - r_A| \leq K|t_B)} + p \sum_{t_A \geq t_f} \frac{u(r_A|t_A)u(t_A)u(r_B)}{u(r_A)u(|r_B - r_A| \leq K)}.$$

Proof: The proof can be found in the Appendix. ■

Utility Transitivity: Timer values have the desirable behavior that their expected value increases as a function of distance. However, timer values quickly become poorer indicators of proximity as their value increases (this is referred to as the “distance effect” [25]). In order to improve the efficiency of utility-based routing it is therefore desirable to reduce the uncertainty for higher timer values. This can be achieved by transitivity.

Let a node A with timer value $\tau_A(D)$ encounter another node B with timer value $\tau_B(D)$, such that $\tau_B(D) \ll \tau_A(D)$. Let also d_{AB} denote the physical distance between A and B , and let $\bar{t}_{mm}(d_{AB})$ denote the expected time it takes a node to move a distance d_{AB} , under mobility model “mm”. Then, on average $\tau_A(D) - \tau_B(D)$ should not exceed $\bar{t}_{mm}(d_{AB})$. If it does, then the information contained in the two timers $(\tau_A(D), \tau_B(D))$ contradicts the position information (d_{AB}) . Hence, one of the timers needs to be corrected.




Fig. 1. Probability density function for the timer value far from the destination with and without transitivity (100×100 grid).




Fig. 2. Comparison of analytical and simulation results for upper and lower bounds on delay (left) and randomized routing (right).

For most noncontrived mobility models $\tau_B(D) \ll \tau_A(D)$ implies that $d_{BD} < d_{AD}$ (with high probability). Thus, it is significantly more probable that it's the larger timer value that needs to be adjusted according to the smaller timer value. A reasonable choice for this adjustment would be $\bar{t}_{mm}(d_{AB})$, since $E[\tau_A(D) - \tau_B(D)] = \Theta(\bar{t}_{mm}(d_{AB}))$. If a node moves according to the random waypoint model, it takes approximately d_{AB} steps on average to move a distance of d_{AB} (assuming d_{AB} is smaller than an epoch). Similarly, in the case of random walks, it takes a node d_{AB}^2 on average to move a distance of d_{AB} [35]. Summarizing, the transitivity functions for the mobility models we use here are given by:

(Random Waypoint) if $\tau_B(D) < \tau_A(D) - d_{AB}$, set $\tau_A(D) = \tau_B(D) + d_{AB}$;

(Random Walk) if $\tau_B(D) < \tau_A(D) - d_{AB}^2$, set $\tau_A(D) = \tau_B(D) + d_{AB}^2$.

As an example of the beneficial effect of transitivity on the statistics of the utility function, see Fig. 1, where a significant decrease in variance is achieved.

E. Simulation versus Analysis

In this final section, we compare our analytical results, regarding the expected delay of different algorithms, to simulation results. We first look into the upper and lower bounds on the expected delay achievable by any scheme, namely the expected delay of Direct Transmission and of the Optimal scheme, respectively. (Note that these are bounds on the delay of both single and multi-copy schemes.) In the left plot of Fig. 2, we fix M to 20, K to 5, and compare simulation and analytical results for the expected delay of these two schemes, as a function of the torus size N . As one can see from this figure, simulation and analytical plots present a relatively close match for both Direct Transmission (some minor discrepancies are due to some approximations, stated in the proofs) and the Optimal scheme (which implies that our approximation of exponential meeting times is fairly accurate).




Fig. 3. Comparison of analytical and simulation results for utility-based routing (left) and seek and focus (right) as a function of K .

In the right plot of Fig. 2, we fix N to 2500 and M to 20, and compare analytical and simulation results for Randomized Routing, for increasing transmission range K . We also include analytical and simulation results for the expected delay of direct transmission, and optimal schemes. Note that, in the case of the randomized algorithm, the analytical plot is only an upper bound (as explained in Section IV-C). As can be seen in this figure the efficiency of Randomized Routing increases quickly for higher transmission ranges.

Finally, in Fig. 3, we compare our analytical results for Utility-based Routing and Seek and Focus routing to simulation results for a 50×50 grid with $M = 40$ nodes. (Results correspond to the case of no transitivity for both schemes, and no re-seek phase for Seek and Focus, and were calculated according to Theorem 4.4 and Lemma 4.5). As is evident by this figure, these simulation and analytical results are also closely matched. Furthermore, like Randomized Routing, the efficiency of both schemes increases with transmission range K . However, note that in this small network, utility-based routing does not suffer from the slow-start phase, and thus seek and focus (without the re-seek phase) does not offer considerable improvement.

V. SIMULATION RESULTS

We have also used a custom discrete event-driven simulator to evaluate the performance of different routing protocols under a large range of connectivity levels, under more realistic conditions (e.g., media access, contention), and for other mobility models (e.g., random waypoint). Although the networks of interest are disconnected in general, they may range from extremely sparse to *almost* connected networks, depending on the application. Connectivity is directly related to the number and quality of forwarding opportunities presented to the message custodian, and thus is expected to have an important effect on performance. A slotted CSMA MAC protocol has been implemented, in order to access the shared channel. Finally, we have implemented and simulated the following single-copy protocols (individual protocol parameters were tuned to achieve a good transmissions–delay tradeoff per scenario): (i) randomized routing; (ii) utility-based routing with no transitivity; (iii) utility-based routing with transitivity; (iv) Seek and Focus with transitivity; (v) optimal oracle-based algorithm.

Before we proceed, it is necessary to define a meaningful connectivity metric. There are two types of connectivity that are important in ICMNs, namely “static” and “dynamic” connectivity. Static connectivity indicates how connected a random snapshot




Fig. 4. Number of transmissions and delivery delay for random walk mobility.

of the connectivity graph will be. Although a number of different metrics have been proposed (for example [44]), no widespread agreement exists, especially if one needs to capture both disconnected and connected networks. We believe that a meaningful metric for the networks of interest is the *expected maximum cluster size* defined as the percentage of total nodes in the largest connected component (cluster).¹⁰ We vary the transmission range K in order to span the entire static connectivity range.

Dynamic connectivity can be seen as a measure of how many new nodes are encountered by a given node within some time interval, and is important in situations where mobility is exploited to deliver traffic end-to-end. If nodes move in an IID manner, this is directly tied to the mixing time for the graph representing the network [35]. The larger the mixing time, the more “localized” the node movement, and the longer it will take a node to *carry* a message to a remote part of the network. We first present results for the random walk and random waypoint models. The random waypoint has one of the fastest mixing times ($\Theta(\sqrt{N})$), while the random walk has one of the slowest ($\Theta(N)$) [35].

In all scenarios there are 100 nodes moving inside a 500×500 network. Also, 100 node pairs are chosen randomly and 1–2 packets are sent per pair throughout a run. Each node has a buffer space of 50 messages. (Simulations we performed showed that single-copy schemes are much less sensitive to traffic load and buffer space than multi-copy ones [12].) Finally, all results are averages over 1000 runs.

Random Walk: Fig. 4 depicts the average number of transmissions (per message) and the delivery delay, respectively, when all nodes perform independent random walks. As can be seen there, randomized routing is quite fast for high connectivity levels. The message can quickly search large clusters of nodes, and the penalty of a wrong decision is small since nodes move very slowly. However, this speed comes with 1 (high connectivity) to 2 (low connectivity) orders of magnitude more transmissions than other schemes. On the other hand, utility-based routing without transitivity performs extremely few transmissions, but is extremely slow. Transitivity, by successfully taking advantage of the high locality of random walk movement, can help reduce delay by almost an order of magnitude with essentially no increase in transmissions. Nevertheless, delivery delay is still larger than randomized, especially for high connectivity levels. Seek and Focus combines the advantages of the two schemes efficiently, achieving the small delays of randomized routing with only a small increase in transmissions compared to utility routing. As a final note, we found that, in both this and the next scenarios, the standard deviation of the delay

¹⁰This connectivity metric would be a lower bound on the percentage of total node pairs that are connected (another popular connectivity metric), since it only takes into account the pairs that belong in the largest cluster.




Fig. 5. Number of transmissions and delivery delay for random waypoint mobility.




Fig. 6. Number of transmissions and delivery delay for community-based mobility.

for most schemes was about the same as their average, which is in accordance with our assumption about the exponential nature of quantities involved.

Random Waypoint: Results for the case when nodes move according to the random waypoint model are depicted in Fig. 5. The story here is somewhat different. First, because nodes move quickly around the network, static connectivity does not have as large an effect on performance. Second, randomized routing performs rather poorly here. Because, nodes move fast, the penalty for forwarding the message to the “wrong” node is high. This is also evident by the fact that Seek and Focus does not improve the delay of Utility-based routing with transitivity as much as in the random walk case. Due to increased node movement, the message does not get stuck in local maxima for a long time. Third, it is evident that, despite the higher speed by which location information becomes obsolete, a carefully chosen utility function can help improve performance significantly. Finally, as one can see in both sets of plots, higher connectivity creates more forwarding opportunities, allowing good protocols to perform even better.

Community-Based Mobility: Popular mobility models like the ones we’ve examined so far, assume that each node may move equally frequently to every network location. Furthermore, such models usually assume that all nodes have the same mobility characteristics, that is, every node’s mobility process is identical, and independently distributed. However, numerous recent studies based on mobility traces from real networks (e.g., university campuses, conferences, etc.) have demonstrated that these two assumptions rarely hold in real-life situations [4], [45]. For this reason, we would also like to compare the performance of all protocols under a more realistic mobility model, called “Community-based Mobility Model”, that is motivated by such traces and better resembles real node movement [36].

In Fig. 6, we depict the total transmissions and average delivery delay for the Community-based model with heterogeneous node mobility. The conclusions that can be drawn from these plots are similar to those in the previous two scenarios. Here, because of the high locality of many nodes, timers often contain quite accurate information, and utility-based schemes perform considerably better than randomized routing.

Specifically, both Utility routing with transitivity and Seek and Focus are an order of magnitude faster than Randomized routing. Seek and Focus can actually take advantage of few randomized transmissions and reduce delay a bit further, but only for more connected networks, where larger clusters of nodes could be searched fast. Summarizing, even in this more realistic scenario, our theory provides adequate insight into the performance of different schemes.

VI. CONCLUSION

In this work, we have dealt with the problem of single-copy routing in intermittently connected mobile networks. We have presented a number of increasingly sophisticated single-copy strategies, and used theory and simulations to extensively evaluate their performance. We conclude that carefully designed single-copy schemes can often present competitive alternatives to more general multi-copy strategies, without the overhead of using redundant copies. Furthermore, we have argued that the understanding provided by such a detailed examination can be valuable for designing better, general-purpose multi-copy schemes. As evidence of this, we have used some of the schemes presented here to create a multi-copy scheme in [14] that can achieve close to optimal performance in a wide range of scenarios.

In future work, we plan to investigate more general utility functions, that can capture and take advantage of a wider range of mobility characteristics pertinent to more realistic mobility models, that might exhibit correlation in both space and time.

APPENDIX

Proof of Lemma 4.1: The electrical network analogy for random walks on graphs is the following [35], [40]: Each vertex of the graph corresponds to a node in an electrical network; each edge corresponds to a resistance of 1 Ohm; instead of using involved combinatorics to count the number of a set of paths and its respective probability for a random walk, one can calculate effective resistances between given points in the corresponding electric network.

For example, let d denote the average node degree of the graph; then it is known [35], [40] that

$$P(T_{A(K)} < T_j^+) = 1/dR_{jA}$$

where R_{jA} is the *effective resistance* between j and $A(K)$. In the case of the 2-D torus, $d = 4$, and we can calculate the effective resistance as follows: We inject a current $I = 1$ in j , and draw a current $I = -1$ from $A(K)$ (all nodes of $A(K)$ are short-circuited). We then calculate the voltage difference V_{jA} on the path (i_0, i_1, \dots, i_K) , where

$$i_n = (x_j, y_j + n).$$

R_{jA} is then given by V_{jA}/I . Now, the total number of paths of length K (i.e., shortest paths to $A(K)$) is equal to $4(2^K - 1)$. Furthermore, the total number of such paths containing edge i_n, i_{n+1} is equal to $2^{K-n} - 1$, which implies that the total current

passing through this edge is equal to $(2^{K-n} - 1)/4(2^K - 1)$. Therefore, R_{jA} is given by

$$R_{jA} = \sum_{n=0}^{K-1} \left(\frac{2^{K-n} - 1}{4(2^K - 1)} \right) = \frac{2^{K+1} - K - 2}{4(2^K - 1)}.$$

Replacing d and R_{jA} in $P(T_{A(K)} < T_j^+) = 1/dR_{jA}$, we get (1).

Proof of Lemma 4.3:

- i) Since the node may lie on any position of the torus with equal probability $1/N$, we can easily calculate $u(r)$ by counting the number of positions at a given distance r from the origin.
- ii) We know that the Markov Chain corresponding to the random walk on the 2-D torus is reversible (see [35]). Therefore, the probability distribution of the time elapsed since the walk *last* encountered the destination (i.e., $u(t)$) is equal to the probability distribution of the time until the walk *first* hits the destination, starting from the stationary distribution. Theorem 4.1 tell us that this time is exponentially distributed with average ET_{RW} .
- iii) $u(r|t)$ gives us the probability distribution of the walk's position given that it has not seen the destination for t time units. In other words, it has not visited any state at distance smaller than $K + 1$ from the destination. Therefore, its past movement can be described by the transition matrix U , where the transition probabilities from states (positions) at distance $K + 1$ to states at distance K are set to zero, and the respective probabilities are equally distributed to all other allowable moves from a state at distance $K + 1$. It is then easy to see that $u(r|t)$ will be given by the t -step transition probabilities U^t with an initial distribution that is uniform right on the border outside of the "in range" area, captured by u_0 .

Proof of Lemma 4.4: We know that A is at distance r_A . A 's timer may have any value t_A with probability $u(t_A|r_A) = u(r_A|t_A)u(t_A)/u(r_A)$ (Bayes rule). We also know that B 's timer probability distribution is $u(t)$ and we're interested in the cases where $t_B < t_A$ (since only then is the message forwarded). Finally, we know that B is within range of A (i.e., $|r_B - r_A| \leq K$). Therefore, the probability that B is at position r_B , given its timer value t_B and the fact that is within range of A is equal to

$$\frac{u(r_B|t_B)}{u(|r_B - r_A| \leq K|t_B)}$$

where

$$u(|r_B - r_A| \leq K|t_B) = \sum_{r_B:|r_B-r_A|\leq K} u(r_B|t_B).$$

Putting all these together, we get the probability of a transmission jump from state r_A to state r_B . Note that we have assumed that $U_{th} = 0$ (a message gets forwarded as long as the next hop has a smaller timer). However, the procedure can easily be generalized for any utility threshold. One only needs to replace

$t_B < t_A$ in the second summation with $t_B < t_A - t_{th}$, where $t_{th} = U^{-1}(U_{th})$.

Proof of Lemma 4.5: The first term of $p_{tx}(r_A \rightarrow r_B)$ corresponds to the focus phase, where utility-based forwarding is performed. Therefore, we only consider the terms for which $t_A < t_f$. Everything else is the same as in Lemma 4.4. The second term corresponds to the seek phase. If $t_A \geq t_f$ then A does not look into B's timer, but just gives it the message with probability p .

REFERENCES

- [1] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. S. Peh, and D. Rubenstein, "Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebraNet," in *Proc. ACM ASPLOS*, 2002.
- [2] Disruption Tolerant Networking. [Online]. Available: <http://www.darpa.mil/ato/solicit/DTN/>
- [3] H. Wu, R. Fujimoto, R. Guensler, and M. Hunter, "MDDV: Mobility-centric data dissemination algorithm for vehicular networks," in *Proc. ACM SIGCOMM Workshop on Vehicular Ad Hoc Networks (VANET)*, 2004.
- [4] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [5] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Proc. ACM SIGCOMM*, 2004.
- [6] R. C. Shah, S. Roy, S. Jain, and W. Brunette, "Data mules: Modeling and analysis of a three-tier architecture for sparse sensor networks," *Elsevier Ad Hoc Netw. J.*, 2003.
- [7] Delay Tolerant Networking Research Group. 2007 [Online]. Available: <http://www.dtnrg.org>
- [8] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *Proc. Mobile Computing and Networking*, 1998.
- [9] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Trans. Networking*, vol. 10, no. 4, pp. 477–486, Aug. 2002.
- [10] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic routing in intermittently connected networks," *SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 7, no. 3, 2003.
- [11] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Duke Univ., Durham, NC, Tech. Rep. CS-200006, Apr. 2000.
- [12] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and wait: Efficient routing in intermittently connected mobile networks," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [13] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu, "The Broadcast Storm problem in a mobile ad hoc network," *Wireless Networks*, vol. 8, no. 2/3, 2002.
- [14] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: The multiple-copy case," *IEEE Trans. Networking*, vol. 16, no. 1, Feb. 2008.
- [15] M. Mauve, J. Widmer, and H. Hartenstein, "A survey on position-based routing in mobile ad hoc networks," *IEEE Network*, vol. 15, no. 6, pp. 30–39, Nov.-Dec. 2001.
- [16] N. Sadagopan, F. Bai, B. Krishnamachari, and A. Helmy, "Paths: Analysis of path duration statistics and their impact on reactive manet routing protocols," in *Proc. ACM MOBIHOC*, 2003.
- [17] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *Proc. ACM MOBIHOC*, 2004.
- [18] Q. Li and D. Rus, "Communication in disconnected ad hoc networks using message relay," *J. Parallel Distrib. Comput.*, vol. 63, no. 1, pp. 75–86, 2003.
- [19] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: An approach to interplanetary Internet," *IEEE Commun. Mag.*, vol. 41, no. 6, pp. 128–136, Jun. 2003.
- [20] E. P. C. Jones, L. Li, and P. A. S. Ward, "Practical routing in delay-tolerant networks," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [21] T. Small and Z. Haas, "Resource and performance tradeoffs in delay-tolerant wireless networks," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.

- [22] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," in *Proc. IFIP Networking*, 2006.
- [23] J. Widmer and J.-Y. L. Boudec, "Network coding for efficient communication in extreme networks," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [24] J. Leguay, T. Friedman, and V. Conan, "DTN routing in a mobility pattern space," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [25] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli, "Age matters: Efficient route discovery in mobile ad hoc networks using encounter ages," in *Proc. ACM MOBIHOC*, 2003.
- [26] Y. Wang, S. Jain, M. Martonosi, and K. Fall, "Erasure coding based routing for opportunistic networks," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [27] G. Sharma and R. Mazumdar, "On achievable delay/capacity trade-offs in mobile ad hoc networks," in *Proc. IEEE WiOpt*, 2004.
- [28] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah, "Throughput-delay trade-off in wireless networks," in *Proc. IEEE INFOCOM*, 2004, pp. 464–475.
- [29] M. Neely and E. Modiano, "Capacity and delay tradeoffs for ad hoc mobile networks," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 1917–1937, Jun. 2005.
- [30] R. M. de Moraes, H. R. Sadjadpour, and J. Garcia-Luna-Aceves, "Throughput-delay analysis of mobile ad hoc networks with a multi-copy relaying strategy," in *Proc. IEEE Conf. Sensor and Ad Hoc Communications and Networks (SECON)*, 2004, pp. 200–209.
- [31] S. Ioannidis and P. Marbach, "Towards an understanding of EASE and its properties," in *Proc. IEEE WiOpt*, 2004, pp. 267–276.
- [32] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Single-copy routing in intermittently connected mobile networks," in *Proc. IEEE Conf. Sensor and Ad Hoc Communications and Networks (SECON)*, 2004, pp. 235–244.
- [33] *Specification of the Bluetooth System 1.2*, The Bluetooth Special Interest Group, 2004.
- [34] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient Routing in intermittently connected mobile networks: the single-copy case," Univ. Southern California, Los Angeles, Tech. Rep. CENG-2005-10, 2005.
- [35] D. Aldous and J. Fill, "Reversible Markov Chains and random walks on graphs." (Monograph in preparation.) [Online]. Available: <http://stat-www.berkeley.edu/users/aldous/RWG/book.html>
- [36] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Performance analysis of mobility-assisted routing," in *Proc. ACM/IEEE MOBIHOC*, 2006.
- [37] B. Karp and H. T. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proc. ACM MOBIHOC*, 2000.
- [38] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications*. Boston, MA: Birkhauser, 1998.
- [39] A. Jindal and K. Psounis, "Performance analysis of epidemic routing under contention," in *Proc. ACM IWCMC*, 2006.
- [40] A. K. Chandra, P. Raghavan, W. L. Ruzzo, and R. Smolensky, "The electrical resistance of a graph captures its commute and cover times," in *Proc. ACM Symp. Theory of Computing (STOC)*, 1989.
- [41] Torus Hitting Times and Green's Functions. [Online]. Available: <http://www.math.tamu.edu/rellis/comb/torus/torus.html>
- [42] J. Y. Le Boudec and M. Vojnovic, "Perfect simulation and stationarity of a class of mobility models," in *Proc. IEEE INFOCOM*, 2005, pp. 2743–2754.
- [43] R. Durrett, *Probability: Theory and Examples*, 2nd ed. London, U.K.: Duxbury Press, 1995.
- [44] O. Dousse, P. Thiran, and M. Hasler, "Connectivity in ad hoc and hybrid networks," in *Proc. IEEE INFOCOM*, 2002, pp. 1079–1088.
- [45] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," *ACM Mobile Comput. Commun. Rev.*, vol. 9, no. 2, pp. 40–55, Apr. 2005.



Thrasyvoulos Spyropoulos (S'02) was born in Athens, Greece, in July 1976. He received the Diploma in electrical and computer engineering from the National Technical University of Athens, Greece, in February 2000, and the Ph.D. degree in electrical engineering from the University of Southern California, Los Angeles, in May 2006.

He was a Post-Doctoral Researcher at INRIA, Sophia Antipolis, France. He is currently a Senior Researcher at ETH, Zurich, Switzerland.

Dr. Spyropoulos has been a member of the Association for Computing Machinery (ACM) since 2006.



Konstantinos Psounis (M'08) was born in Athens, Greece, in November 1974. He received the Diploma in electrical and computer engineering from the National Technical University of Athens, Greece, in June 1997, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in December 2002.

He is an Assistant Professor of electrical engineering and computer science at the University of Southern California, Los Angeles.

Dr. Psounis has been a member of the Association for Computing Machinery (ACM) since 2001.



Cauligi S. Raghavendra (F'07) received the B.Sc. (Hons.) physics degree from Bangalore University, India, in 1973, the B.E. and M.E. degrees in electronics and communication from the Indian Institute of Science, Bangalore, in 1976 and 1978, respectively, and the Ph.D. degree in computer science from the University of California at Los Angeles in 1982.

He is a Professor of electrical engineering and computer science and Senior Associate Dean for Strategic Initiatives of the Viterbi School of Engineering at the University of Southern California, Los Angeles.

Dr. Raghavendra was a recipient of the Presidential Young Investigator Award for 1985.