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

Thrasyvoulos 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, etc. In this context, conventional routing schemes 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. Furthermore, proposed efforts to reduce the overhead of flooding-based schemes have often been plagued by large delays. With this in mind, we introduce a new family of routing schemes that "spray" a few message copies into the network, and then route each copy independently towards the destination. We show that, if carefully designed, spray routing not only performs significantly fewer transmissions per message, but also has lower average delivery delays than existing schemes; furthermore, it is highly scalable and retains good performance under a large range of scenarios.

Finally, we use our theoretical framework proposed in our 2004 paper to analyze the performance of spray routing. We also use this theory to show how to choose the number of copies to be sprayed and how to optimally distribute these copies to relays.

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

I. Introduction

WIRELESS data networks often aim at extending Internet services into the wireless domain. Services like GPRS enable Internet access through the widespread cellular infrastructure, while the deployment of WiFi 802.11 access points provides direct Internet connectivity for wireless users (mainly laptops and PDAs) that are within range. Additionally, self-organized ("ad hoc" or "peer-to-peer") wireless networks have been proposed for applications where setting up a supporting, wired infrastructure might be too costly (e.g., sensor networks) or simply not an option (e.g., disaster relief, deep space networks).

Despite these ongoing efforts, wireless access currently seems to give rise to inconvenience and frustration more often than providing the envisioned flexibility to the user. Cellular access is low bandwidth and expensive, while WiFi access is typically only available at a few "hotspots" that the user has to locate and move to, without real "mobile computing". Further, ad hoc networks have yet to find much application outside the research or military community, while some dire issues regarding their scalability properties have been identified [2].

The reason for these failures is that many of the assumptions made in the wired world, and which are largely responsible for the success of the Internet, do not hold in the wireless environment. The concept of a connected, stable network over which data can be routed reliably rarely holds there. Wireless signals are subject to multi-path propagation, fading, and interference making wireless links unstable and lossy. Additionally, frequent node mobility (e.g., as in vehicular ad hoc networks—VANETs [3]) significantly reduces the time a "good" link exists, and constantly changes the network connectivity graph. As a result, wireless connectivity is volatile and usually intermittent, as nodes move in and out of range from access points or from each other, and as signal quality fluctuates.

In addition to the cases of wireless Internet access and ad hoc networks, the need to depart from the traditional networking practices has been recognized for a number of emerging wireless applications. Sensor networks can significantly increase their lifetime by powering down nodes often, or by using very low power radios. This implies that many links will be down frequently, and complete end-to-end paths often will not exist [4]. Tactical networks may also choose to operate in an intermittent fashion for LPI/LPD reasons (low probability of interception and low probability of detection) [5]. Finally, deep space networks [6] and underwater networks [7] often have to deal with long propagation delays and/or intermittent connectivity, as well. These new networks are often referred to collectively as Delay Tolerant Networks (DTN [8]). What they all share in common is that *they can neither make any assumptions about the existence of a contemporaneous path to the destination nor assume accurate knowledge of the destination's location or even address, beforehand.*

Under such intermittent connectivity many traditional protocols fail (e.g., TCP, DNS, etc. [8], [9]). It is for this reason that novel networking architectures are being pursued that could provide mobile nodes with better service under such intermittent characteristics [9], [10]. Arguably though, the biggest challenge to enable networking in intermittently connected environments is that of routing. Conventional Internet routing protocols (e.g., RIP and OSPF), as well as routing schemes for mobile ad hoc networks such as DSR, AODV, etc. [11], assume that a complete path exists between a source and a destination, and try to

Manuscript received August 29, 2005; revised June 25, 2006, and December 7, 2006; approved by IEEE/ACM Transactions on Networking Editor S. Palazzo. 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.2007.897964

discover these paths *before any useful data is sent*. Thus, if no end-to-end paths exist most of the time, these protocols fail to deliver any data to all but the few connected nodes.

However, this does not mean that packets can never be delivered in these networks. Over time, different links come up and down due to node mobility. 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 (e.g., a new node moves in range or an existing one wakes-up), and so on and so forth, until it reaches its destination. This model of routing constitutes a significant departure from existing routing practices. It is usually referred to as "mobility-assisted" routing, because node mobility often needs to be exploited to deliver a message to its destination (other names include "encounter-based forwarding" or "store-carry-and-forward"). Routing here consists of independent, local forwarding decisions, based on current connectivity information and predictions of future connectivity information, and made in an *opportunistic* fashion. The crucial question any routing algorithm has to answer in this context is "who makes a good next hop when no path to the destination currently exists and/or no other information about this destination might be available?".

Despite a number of existing proposals for opportunistic routing [12]–[16] the answer to the previous question has usually been "everyone" or "almost everyone". The majority of existing protocols are flooding-based that distribute duplicate copies to all nodes in the network [12] or a subset of them (e.g., gossiping [15], and utility-based flooding [14]). Although flooding can be quite fast in some scenarios, the overhead involved in terms of bandwidth, buffer space, and energy dissipation is often prohibitive for small wireless devices (e.g., sensors). We call schemes like these, which use more than one copy per message, "multi-copy" schemes. Single-copy schemes that only route one copy per message can considerably reduce resource waste [1], [4]. Yet, they can often be orders of magnitude slower than multi-copy algorithms and are inherently less reliable. These latter characteristics might make single-copy schemes very undesirable for some applications (e.g., in disaster recovery networks or tactical networks beyond enemy lines; even if communication *must* be intermittent, minimizing delay or message loss is a priority). Summarizing, *no routing scheme for intermittently connected environments currently exists that can achieve both small delays and prudent usage of the network and node resources.*

For this reason, we propose a family of multi-copy protocols called *Spray routing*, which can achieve both good delays and low transmissions. Spray routing algorithms generate only a small, carefully chosen number of copies to ensure that the total number of transmissions is small and *controlled*. Then, an appropriate single-copy algorithm (based on the insight acquired from our study of single-copy routing strategies [17]) is used to route each of the copies independently. From the perspective of functionality, spray routing can be viewed as a tradeoff between single and multiple copy techniques. Despite this, theory and simulations show that spray routing: (i) achieves an order of magnitude reduction in transmissions compared to flooding-based schemes, and even fewer transmissions than some single-copy schemes; (ii) can at the same time achieve better delays than all existing schemes in most scenarios, if carefully designed; and (iii) has very desirable scalability characteristics, with its relative performance improving as the network size increases.

Additionally, we provide a theoretical analysis of the performance of spray routing that holds for a number of popular and recent mobility models. In addition to the inherent value of such an analysis for performance prediction, we also use our theory to perform system design. Specifically, we provide an efficient algorithm that each node can use to locally choose the number of copies to generate in a given scenario, and also show how to optimally distribute these copies. Our theory also provides a "knob" to explicitly take advantage of different tradeoffs between resource usage and performance, allowing a wireless device to adapt to the individual user's needs, even in situations where a node initially has little or no information about the network.

In the next section, we go over some existing related work. Section III describes our proposed solution, Spray routing, in detail. We deal with the issue of optimal spraying algorithms and analyze their expected delay in Section IV, and then in Section V we show how to optimally choose the total number of copies per message to be sprayed. Simulation results are presented in Section VI, where the performance of Spray routing is compared against a number of existing schemes with respect to the message delivery delay and the total number of transmissions. Finally, Section VII concludes the paper.

II. Related Work

One approach to deal with very sparse networks or connectivity "disruptions" [5] 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 [18], [19]. In yet other cases, connectivity might be predictable, even though its intermittent (e.g., planetary and satellite movement in Inter-planetary Networks—IPN [6]). Traditional routing algorithms could then be adapted to compute shortest delivery time paths by taking into account future connectivity [20], [21]. 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.

There exists a growing amount of work on opportunistic, DTN routing algorithms. One of the simplest approaches is to let the source or a moving relay node (*DataMule*) carry the message all the way to the destination (*Direct Transmission*) [4]. Although this scheme performs only one transmission, it is extremely slow [22]. Other single-copy schemes have also been explored that can forward a message to improve end-to-end delay [1]. Yet, an even faster way to perform routing in intermittently connected mobile networks (or ICMNs), called *Epidemic Routing*, is to flood the message throughout the network [12]. Although this scheme is guaranteed to find the shortest path when no contention exists for shared resources like wireless bandwidth and buffer space, 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 due to congestion [14], [23], [24].

A number of approaches have been taken to reduce the overhead and improve the performance of epidemic routing [13]–[16], [24], [25]. In [25] the authors examine a number of different strategies to suppress redundant transmissions and clean up valuable buffer space after a message has been delivered with epidemic routing. In [15], [24] a message is forwarded to another node with some probability smaller than one (i.e., data is "gossiped" instead of flooded). Finally, in [13] a simple method to take advantage of the history of past encounters is implemented in order to make fewer and more "informed" forwarding decisions than epidemic routing. The concept of history-based or utility-based routing is further elaborated in [14], [26] and has also been studied for regular, connected networks [27]. Results from these works indicate that using the age of last encounter with a node, when making a forwarding decision, results in superior performance than flooding. Finally, it has also been proposed that ideas from the area of *Network Coding* could be useful to reduce the number of bytes transmitted by flooding [16]. Despite the large number of existing approaches, most proposed schemes are based on epidemic-routing or some other form of controlled flooding [12]–[14], and, thus, are plagued by the shortcomings of flooding-based schemes [23], [24].

The idea of "spraying" is also not entirely new. For example, in cellular networks it has been used to deliver data to nodes that are highly mobile and change their attachment point frequently [28]. Instead of sending the message only to the base station where the node was last seen, duplicate messages are also sent to other, nearby base stations. Furthermore, the idea of using a few redundant copies to improve delay has also been proposed elsewhere in both connected and disconnected ad hoc networks [25], [29], [30]. However, these works do not answer important questions like how the copies should be distributed, or how many of them. An interesting effort to address such issues, yet in the context of reliability, is undertaken in [31]. Finally, a growing number of efforts has been devoted to collect real mobility traces [32]–[34] and use them to evaluate the feasibility of the DTN approach [34].

Also, 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 [22], and of related schemes exploiting mobility [29], [30], [35], [36]. 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, where all related analytical results are strictly a function of the number of nodes (other parameters scale to ensure connectivity) [29], [35]. 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 we would like to analyze.

In the context of disconnected networks, the performance of epidemic routing with or without contention has been addressed in a number of works [1], [15], [37]–[39]. Additionally, there also exist some efforts to analyze the performance of 2-hop schemes using redundant copies [25], [39]. However, these works often ignore the effect of different copy distribution policies, and use simulations to acquire some of the model parameters [25], [29], [39]. This significantly reduces the usefulness of analytical expressions. In this paper, we address all these issues and provide analytical expressions for the performance of spraying under a range of mobility models.

III. Spray Routing

In this section, we explore the problem of efficient routing in intermittently connected mobile networks (ICMNs), and describe our proposed solution, Spray routing. Our problem setup consists of a number of nodes moving inside a bounded area according to a stochastic mobility model. Additionally, we assume that the network is disconnected at most times, and that transmissions are faster than node movement (i.e., it takes less time to transmit a message using the wireless medium—ignoring queueing delay—than to move it physically for the same distance using node mobility[1]).

Our study of single-copy routing algorithms [17] showed that using only one copy per message is often not enough to deliver a message with high reliability and relatively small delay. At the same time, routing too many copies in parallel, as in the case of epidemic routing or gossiping, can often have disastrous effects on performance (as is evident also from Fig. 2 later in Section IV-C). In addition to the very high number of transmissions, flooding-based schemes begin to suffer severely from contention as traffic increases, and their delay increases rapidly. Based on these observations, we have identified the following desirable design goals for a routing protocol in intermittently connected mobile networks:

- perform significantly fewer transmissions than flooding-based routing schemes, under all conditions.
- deliver a message faster than existing single and multi-copy schemes, and exhibit close to optimal delays.
- deliver the majority of the messages generated;

Additionally, we would like this protocol to also be:

- highly scalable, that is, maintain the above performance behavior despite changes in network size or node density.
- simple, and require as little knowledge about the network as possible, in order to facilitate its implementation.

A. "Spray and Wait" Routing

Since too many transmissions are detrimental on performance, especially as the network size increases, our first protocol, *Spray and Wait*, distributes only a small number of copies each to a different relay. Each copy is then "carried" all the way to the destination by the designated relay.

Definition 3.1 (Spray and Wait): Spray and Wait routing consists of the following two phases:

[1]This is reasonable assumption with modern wireless devices. Assume, 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.

- *spray phase*: for every message originating at a source node, L message copies are initially spread—forwarded by the source and possibly other nodes receiving a copy—to L distinct relays. (Details about different spraying methods will be given later.)
- *wait phase*: if the destination is not found in the spraying phase, each of the L nodes carrying a message copy performs "Direct Transmission" (i.e., will forward the message only to its destination).

Spray and Wait decouples the number of transmissions per message from the total number of nodes. Thus, transmissions can be kept small and essentially fixed for a large range of scenarios. Additionally, its mechanism combines the speed of epidemic routing with the simplicity and thriftiness of direct transmission. Initially, it "jump-starts" spreading message copies quickly in a manner similar to epidemic routing. However, it stops when enough copies have been sprayed to guarantee that at least one of them will reach the destination, with high probability. If nodes move quickly enough around the network or "cover" a sizeable part of the network area in a given trip, we will show that *only a small number of copies can create enough diversity to achieve close-to-optimal delays*. Some examples of applications with such favorable mobility characteristics would be Vehicular Ad hoc Networks for real-time traffic reports and accident prevention [3], or a wireless mesh network over city buses equipped with radios [20].

As we mentioned earlier, the basic idea behind Spray and Wait (i.e., extending the 2-hop scheme of [22] to introduce more than one relays) is relatively simple and has been identified as beneficial by other researchers also [25], [29], [30]. However, a number of important questions need to be answered first, before the desirable performance can be achieved: (i) How many message copies should one use in a given scenario? (ii) How should these copies be optimally distributed? (iii) How small delays can one achieve in various settings, without compromising the requirement for very few transmissions? and (iv) How does the performance of Spray and Wait scale, as the number of nodes in the network increases to the limit? We will be answering each of these questions in the remaining of the paper.

B. *"Spray and Focus" Routing*

Although Spray and Wait combines simplicity and efficiency, there are some situations were it might fall short. As explained earlier, it requires the existence of enough nodes that roam around the network often, which could potentially carry a message to a destination that lies far. Usually, Spray and Wait spreads all its copies quickly to the node's immediate neighborhood. Hence, if the mobility of each node is restricted to a small local area, then none of the nodes carrying a copy might ever see the destination.

An example where such localized mobility might arise could be, for example, a university campus, where most people tend to stay or move locally within their buildings for long stretches of time [33]. In such situations, partial paths may exist over which a message copy could be quickly transmitted closer to the destination. Yet, in Spray and Wait a relay with a copy will naively wait until it moves within range of the destination itself. This problem could be solved if some other single-copy scheme is used to route a copy after it is handed over to a relay, a scheme that takes advantage of transmissions (unlike Direct Transmission).

We propose the use of the single-copy utility-based scheme from [17] for this purpose. Each node maintains a timer for every other node in the network, which records the time elapsed since the two nodes last encountered each other[2] (i.e., came within transmission range). These timers are similar to the *age of last encounter* in [27], and are useful, because they contain indirect (relative) location information. Specifically, for a large number of mobility models, it can be shown that a smaller timer value on average implies a smaller distance from the node in question. Further, we use a "transitivity function" for timer values (see details in [17]), in order to diffuse this indirect location information much faster than regular last encounter based schemes [27]. The basic intuition behind this is the following: in most situations, if node B has a small timer value for node D, and another node A (with no info about D) encounters node B, then A could safely assume that it is also probably close to node D. We assume that these timers are the *only* information available to a node regarding the network (i.e., no location info, etc.).

We have seen in [17] that appropriately designed utility-based schemes, based on these timer values, have very good performance in scenarios were mobility is low and localized. This is the exact situation were Spray and Wait loses its performance advantage. Therefore, we propose a scheme were a fixed number of copies are spread initially exactly as in Spray and Wait, but then each copy is routed independently according to the *single-copy* utility-based scheme with transitivity [17]. We call our second scheme *Spray and Focus*.

Definition 3.2 (Spray and Focus): Spray and Focus routing consists of the following two phases:

- *spray phase*: for every message originating at a source node, L message copies are initially spread—forwarded by the source and possibly other nodes receiving a copy—to L distinct "relays".
- *focus phase*: let $U_X(Y)$ denote the utility of node X for destination Y; a node A, carrying a copy for destination D, forwards its copy to a new node B it encounters, if and only if $U_B(D) > U_A(D) + U_{th}$, where U_{th} (utility threshold) is a parameter of the algorithm.

IV. Performance of Spray Routing

In this section, we will analyze the delay of Spray routing. In addition to the intrinsic value of such a theoretical analysis, which is the ability to predict the performance of the schemes in a larger range of scenarios than simulations or experiments can, we also need this theory to do system design. First, we would like to know what is the right number of copies to be sprayed, in order to achieve good performance for Spraying algorithms. Without this number spraying performance could be as bad as that of Direct Transmissions or Epidemic routing in different scenarios. Second, we envision many situations where the user

[2]In practical situations, each node would actually maintain a cache of the most recent nodes that it has encountered, in order to reduce the overhead involved in a large network.

(or his equipment on his behalf) should be able to evaluate the potential performance benefits by using some extra copies, and compare them with the respective costs (e.g., potential energy costs, "forwarding credits" spent, or even monetary costs) so as to decide on the specific tradeoffs that suit him personally. However, in order to be able to do so, we need to first derive an appropriate delay expression that quantifies the effect of using a given number of copies on performance.

Finally, if mobility-assisted routing and delay tolerant networking is to ever become acceptable by the user as a *useful* alternative for mobile networking, it will be necessary to provide him with some sort of (frequently updated) estimates or predictions of the level of service that he should expect (e.g., "currently connected to Internet: all services available" or "frequent disconnections: e-mail and limited web access available"). In a connected environment, direct measurements (e.g., round-trip time) would serve this purpose. However, in the intermittent environment, collecting such end-to-end statistics is much more difficult. Instead, local measurements should be used to predict end-to-end values based on some appropriate theoretical model.

Throughout this and the next section, we will be making the following assumptions:

Network: M nodes move on a $\sqrt{N} \times \sqrt{N}$ two-dimensional torus. Each node can transmit up to distance $K \geq 0$ meters away, where K is much smaller than the value required for connectivity [40]. We assume that links are bi-directional, and that each message transmission takes one time unit.

Mobility Models: We assume that all nodes move according to some stochastic mobility model ("MM"), whose "meeting times" are approximately exponentially distributed or have an exponential tail with expected meeting time equal to EM_{mm} (see [41] for a rigorous definition of *hitting* and *meeting times*). It has been shown that a number of popular mobility models like Random Walk [41], Random Waypoint and Random Direction [25], [42], as well as more realistic, synthetic models based on these [42] exhibit such (approximately) exponential encounter characteristics. Therefore, the analysis and algorithms of this and the following section apply to all these models.

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 flooding-based schemes (as we shall show later in our simulations), we argue that it is significantly less of an issue for our spraying schemes that perform only a handful of transmissions most of the time. Also, in networks that are quite sparse, we expect that only a few nodes would be close enough each time to compete for the same bandwidth. Therefore, we choose to ignore contention, in order to be able to derive useful closed form expressions. We will show that the error introduced by this assumption is small *for the case of spraying schemes*.

A. Optimal Spraying

In both Spray and Wait and Spray and Focus there are L copies that need to be spread initially to L different relays. The first interesting question to be answered is how these L copies should be distributed. (We will talk about how to choose the value of L later.) The choice of spraying method directly affects the expected delay of the spraying phase. Further, this delay is independent of the particular single-copy scheme that is used to route each copy in the second phase.

A number of different spraying heuristics can be envisioned. For example, the simplest way is to have the source node forward all L copies to the first L distinct nodes it encounters (we call this scheme "Source Spraying"). A better way is the following.

Definition 4.1 (Binary Spraying.): The source of a message initially starts with L copies; any node A that has $n > 1$ message copies (source or relay), and encounters another node B (with no copies), hands over to B $\lfloor n/2 \rfloor$ of its copies and keeps $\lceil n/2 \rceil$ for itself; when it is left with only one copy, it switches to Direct Transmission or Utility-based routing, depending on the flavor of spray routing used.

The following theorem states that Binary Spraying is optimal, when node movement is independent and identically distributed (IID).

Theorem 4.1: When all nodes move in an IID manner, Binary Spraying minimizes the expected time until all copies have been distributed.

Proof: Let us call a node "active" when it has more than one copy of a message. Let us further define a spraying algorithm in terms of a function $f : \mathcal{N} \rightarrow \mathcal{N}$ as follows: when an active node with n copies encounters another node, it hands over to it $f(n)$ copies, and keeps the remaining $n - f(n)$. Any spraying algorithm (i.e., any f) can be represented by the following binary tree with the source as its root: assign the root a value of L; if the current node has a value $n > 1$ create a right child with a value of $n - f(n)$ and a left one with a value of $f(n)$; continue until all leaf nodes have a value of 1.

A particular spraying method corresponds then to a sequence of visiting all nodes of the tree. This sequence is random. Nevertheless, *on the average*, all tree nodes at the same level are visited in parallel. Further, since only active nodes may hand over additional copies, the higher the number of active nodes when i copies are spread, the smaller the residual expected delay until all copies are spread. Since the total number of tree nodes is fixed $(2^{1+\log L} - 1)$ for any spraying function f, it is easy to see that the tree structure that has the maximum number of nodes at every level, also has the maximum number of active nodes (on the average) at every step. This tree is the balanced tree, and corresponds to Binary Spraying. ■

B. Delay of Spray Routing

We will first calculate the expected end-to-end delay of our simpler scheme, Spray and Wait. After all copies are distributed, each of the L relays will independently look for the destination (if the latter has not been found yet). In other words, the delay of the wait phase is independent of the spraying method. We compute this delay in the following Lemma.

Lemma 4.1: Let EW denote the expected duration of the "wait" phase, if needed, and let EM_{mm} denote the expected meeting time under the given mobility model. Then, EW is independent of the spraying method used, and given by

$$EW = \frac{EM_{mm}}{L}. \quad (1)$$

Proof: The time until one of the relays finds the destination is the minimum of L independent and exponentially distributed random variables, with average EM_{mm}. ■

Unlike the expected duration of the wait phase, the duration of the spray phase largely depends on the way the L copies are spread. The following theorem calculates the expected delivery time of Binary Spray and Wait. It defines a system of recursive equations that calculates the (expected) residual time after i copies have been spread, in terms of the time until the next copy$(i+1)$ is distributed, plus the remaining time thereafter. It is important to note that the following result is generic. *By plugging into the equations the appropriate meeting time value* EM_{mm}, *we can calculate the expected delay of Spray and Wait for the respective mobility model* [42].

Theorem 4.2: Let $ED_{sw}(L)$ denote the expected delay of the Binary Spray and Wait algorithm, when L copies are spread per message. Let further $ED(i)$ denote the expected remaining delay after i message copies have been spread. Then, $ED(1) \approx ED_{sw}(L)$, where $ED(1)$ can be calculated by the following system of recursive equations:

$$
\begin{aligned}
ED(i) &= \frac{EM_{mm}}{i(M-i)} + \frac{M-i-1}{M-i} ED(i+1), \qquad i \in \left[1, \frac{L}{2}\right];\\
ED(i) &= \frac{ED_{mm}}{i(M-i)} + \frac{M-i-1}{M-i}\\
&\quad \cdot \left(\frac{2i-L}{i} ED(i) + \frac{L-i}{i} ED(i+1)\right),\\
&\quad i \in \left[\frac{L}{2}+1, L-1\right];\\
ED(L) &= EW = \frac{EM_{mm}}{L}.
\end{aligned}
$$

Proof: Let us look into the case, when there are i nodes ($i < L$) that have one or more copies. Further, let us assume that, among the i nodes with copies, X_i of them have more than one (i.e., are "active"), and thus are allowed to forward copies further to other relays. Since all hitting times are independent and exponentially distributed, the time until any of the nodes with a message copy (i) encounters any of the nodes without one $(M-i)$ is equal to $(EM_{mm}/i(M-i))$. Now, if the node encountered is the destination (with probability $1/(M-i)$) the message gets delivered. Otherwise (with probability $(M-i-1)/(M-i)$) the algorithm continues, performing one of the following: a) with probability X_i/i it is one of the active nodes that encountered this other node, and therefore hands it over half its copies; $i+1$ nodes have copies now, and an expected time $ED(i+1)$ remains until delivery; b) with probability $(i-X_i)/i$ it was one of the other nodes carrying a message copy that encountered a new node. Since these relays only forward their message copy to its destination, nothing happens, and the remaining time is still $ED(i)$. Putting it all together

$$
\begin{aligned}
ED(i) &= \frac{EM_{mm}}{i(M-i)}\\
&\quad + \frac{M-i-1}{M-i}\left(\frac{i-X_i}{i} ED(i) + \frac{X_i}{i} ED(i+1)\right).
\end{aligned}
$$

Now, let us represent again a given spraying algorithm as a tree graph (see Theorem 4.1). Let us further look at a given spraying instance, and let us color the tree nodes that have already been reached in the spraying phase. Obviously, the coloring sequence that results in the highest X_i, and therefore the minimum $ED(i)$ for all i, is a breadth-first traversal. In this case, $X_i = i, i \in [1, L/2]$ and $X_i = L-i, i \in [L/2+1, L-1]$. However, a given spraying instance might not exactly follow such a breadth-first traversal. It is possible to derive the exact delay by averaging over all possible traversals, but it involves meticulous calculations that do not offer any interesting insight or significant increase in accuracy (as we shall see in Section IV-C). We therefore assume that breadth-first traversal is always the case, in order to keep our equations simple. Finally, $ED(L) = EW$ is given by (1). ■

The above result, albeit quite useful in accurately predicting the performance of Spray and Wait, is not in closed form. This makes it difficult to theoretically compare the performance of Spray and Wait to that of the optimal scheme [17], or to calculate the number of copies to be used in closed form (more about this in the next Section). For this reason, in the following lemma we also derive an upper bound that is in closed form, by assuming that Source Spray and Wait is performed, that is, only the source can forward a new copy. Note that Source Spray and Wait always has a larger delay than Binary Spray and Wait.

Lemma 4.2: The following upper bound holds for the expected delay of Spray and Wait:

$$ED_{sw} \leq (H_{M-1} - H_{M-L})EM_{mm} + \frac{M-L}{M-1}EW, \quad (2)$$

where H_n is the nth harmonic number, i.e, $H_n = \sum_{i=1}^{n}(1/i) = \Theta(\log n)$.

Proof: Assume that, at some time instant, i of the L copies have already been spread, that is, there are i nodes, including the source node, carrying a message copy. Since only the source can forward another copy, the expected time until another message copy is distributed is equal to the time until the source meets one of the $M-i$ remaining nodes, that is, $EM_{mm}/(M-i)$. Hence, the expected time until $L-1$ different relays are encountered equals $\sum_{i=1}^{L-1}(EM_{mm}/(M-i)) = (H_{M-1} - H_{M-L})EM_{mm}$. This is the time until L message copies, including that of the source, are spread among the M nodes.

Finally, the probability that the destination is not in the first $L-1$ nodes encountered (and, thus, a wait phase is needed) is $1-((L-1)/(M-1)) = (M-L)/(M-1)$. Putting it altogether we get that the expected delay of Spray and Wait is at most

$$\sum_{i=1}^{L-1} \frac{EM_{mm}}{M-i} + \frac{M-L}{M-1}EW.$$

■

This bound becomes pessimistic as the ratio L/M increases. This is because the bound basically includes the full time until all copies are spread, regardless of whether the destination is found in one of the initial steps of the spraying phase. However, when the number of copies is much smaller than the total number of nodes (which is the case of most interest) this bound is tight.

It is important to note that all results regarding the *spraying phase* derived so far for Spray and Wait (i.e., optimal spraying method and its expected delay) hold for Spray and Focus, as well. However, the delay of the two protocols does differ in the second phase. In the focus phase each copy follows a single-copy utility-based algorithm. Utility-based schemes have memory, making it quite more involved to calculate their delay. We have calculated the expected delay for the single-copy utility-based algorithm with a simple utility function in [17], by using a Markov Chain whose state is the distance from the destination. However, we do not know if the distribution of this delay is exponential. In addition, the routing of multiple copies will not be independent, since copies at the beginning of the focus phase would start from roughly the same area of the utility field, and experience similar transmission opportunities (see [17]). Hence, we do not expect something as simple as Lemma 4.1 to hold for Spray and Focus. Nevertheless, we can use the delay of Spray and Wait as an upper bound on the delay of Spray and Focus. We have proven in [1] that utility-based forwarding is (on average) better than not forwarding (i.e., direct transmission). Therefore, the delay of Spray and Focus will be at least as good as that of Spray and Wait, for the same number of copies used (disregarding contention).

As part of future work we plan to investigate the difficult problem of multiple copies being routed in parallel in a given utility field, to see if meaningful theoretical solutions could be drawn. One direction, for example, would be to assume that all copies start from the same position/distance, but then each one is routed independently. This scenario could be modelled with a two-dimensional Markov Chain, whose state is the distance from the destination (as in [17]) *and* the number of copies at that distance. However, we expect these calculations to be quite involved, possibly brute force in the most general case, and thus beyond the scope of this paper.

C. Simulation versus Analysis

In this final section we evaluate the accuracy of our analytical results regarding the expected delay of Spray routing. In Fig. 1, we compare theoretical and simulation results for Binary Spray and Wait, as a function of the number of copies spread L. The simulation plots correspond to "graph simulations" with no contention (only 1 message routed), in order to evaluate the correctness of our analytical expressions. In the left plot, 30 nodes perform independent random walks in a 50×50 network ($K = 0$, i.e., only nodes in the same position can communicate). We also include a plot for the bound of Lemma 4.2, as well as a plot for the optimal delay in this scenario. In the right plot, we compare theoretical and simulation results for Spray and Wait, when 30 nodes perform independent Random Direction mobility in a 500 $\times$ 500 network (pause time is 0, and $K = 30$). As one can see from this figure, simulation and analytical plots for Spray and Wait present a very close match for *both* mobility models. This validates the generality of the analytical expressions (we have also found similar accuracy for Random Waypoint mobility). Additionally, it is evident that the upper bound of (2) is tight for low L/M ratios. Finally, one can see that Spray and Wait can already achieve a delay only 1.5–2 times that of the optimal scheme, using just a handful of transmissions (Epidemic Routing performs $\Theta(M)$ transmission under no contention).

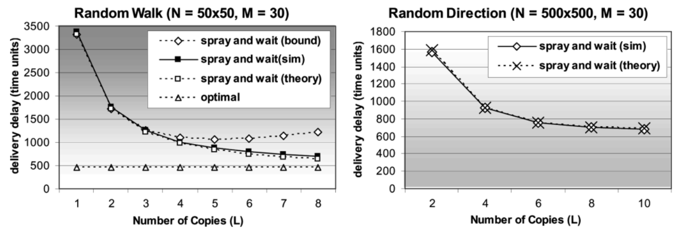

Fig. 1. Comparison of analytical and simulation results for Spray and Wait under Random Walk (left) and Random Direction (right) mobility.

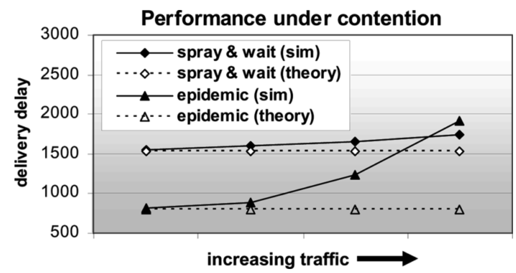

Fig. 2. Comparison of analytical and simulation results for Spray and Wait when there is contention. We assume Random Direction mobility. $N = 500 \times 500, M = 50, K = 20, L = 10$, and average pause time is 100.

We're also interested in how fast our analytical expressions would diverge from real values when increasing amounts of contention (and the overhead of a MAC layer) are introduced (Fig. 2). Let us look first at the plots for Spray and Wait. As explained in the beginning of this section, although contention does affect the accuracy of our theoretical expressions, the error introduced for Spray and Wait is not large (<20%), even for large traffic loads. Therefore, we believe our analytical expressions are useful in predicting performance in more realistic scenarios with contention, as well. We also compare plots for epidemic routing (the theoretical one is for an oracle-based scheme [1]), in order to show how the same traffic loads affect the accuracy of existing analytical expressions for the latter [42]. As is evident by these plots, the actual delays observed for epidemic routing become increasingly worse than what theory predicts. This demonstrates the need to add an appropriate contention model when it comes to modeling flooding-based schemes. A first effort to that direction can be found in [38].

V. Optimizing Spray Routing to Meet Performance Requirements

By definition, many ICMN networks are expected to operate in stressed environments and by nature be *delay tolerant*. Nevertheless, in many situations the network designer or the application itself might still impose certain performance requirements on the protocols (e.g., maximum delay, or maximum transmissions). It is of special interest therefore to examine how Spray routing can be tuned to achieve different levels of performance in a given scenario.

For example, imagine a scenario where a number of users establish a peer-to-peer wireless network for messaging and content sharing [9], [10], and where each node has a limited amount of forwarding requests for other nodes ("credits"). If a user knows how much performance his credit can buy, he has

TABLE I
MINIMUM L TO ACHIEVE EXPECTED DELAY

a	1.5	2	3	4	5	6	7	8	9	10
recursion	21	13	8	6	5	4	3	3	3	2
bound	N.A.	N.A.	11	7	6	5	4	3	3	2
taylor	N.A.	N.A.	10	7	5	4	3	3	3	2

little incentive to spent more (i.e., use more copies) than necessary to achieve his goals.

A. *Choosing L to Achieve a Required Expected Delay*

In this section, we analyze how to choose L (i.e., the number of copies used) in order to achieve a specific expected delay. Let us assume that there is a specific delivery delay constraint to be met. One reasonable way to express such a constraint would be as a factor a times the optimal delay ED_{opt} $(a > 1)$, since this is the best that any routing protocol could do.[3]

Lemma 5.1: The minimum number of copies L_{min} needed for Spray and Wait to achieve an expected delay at most aED_{opt} is independent of the mobility model, the size of the network N, and transmission range K, and only depends on a and the number of nodes M.

The above lemma is straightforward to prove from (2) or Theorem 4.2. The required number of copies $L_{min}(M)$ for Spray and Wait to achieve a desired expected delay can be calculated in any of the following three ways: (i) solve the system of equations of Theorem 4.2 for increasing L, until $ED_{sw}(L) < aED_{opt}$, or (ii) solve the upper bound equation (2) for L, by letting $ED_{sw} = aED_{opt}$, and taking $\lceil L \rceil$, or (iii) approximate the harmonic number H_{M-L} in (2) with its Taylor Series terms up to second order, and solve the resulting third degree polynomial:

$$(H_M^3 - 1.2)L^3 + \left(H_M^2 - \frac{\pi^2}{6}\right)L^2 + \left(a + \frac{2M-1}{M(M-1)}\right)L = \frac{M}{M-1}$$

where $H_n^r = \sum_{i=1}^{n}(1/i^r)$ is the nth harmonic number of order r.

Method (i) is obviously the most accurate one. However, it is also the most cumbersome. Since the upper bound of (2) is tight for small L/M values, if the delay constraint a is not too tight, we can use method (ii) or (iii) to quickly get a good estimate for L_{min}.

In Table I we compare results for L_{min}, as calculated with each of these three methods for different values of a. We assume the number of nodes M equals 100. "N.A" stands for "Non Available" and means that such a low delay value is never achievable by the bound. As can be seen in this table the L found through the approximation is quite accurate when the delay constraint is not too stringent.

B. *Estimating L When Network Parameters are Unknown*

Throughout the previous analysis we have assumed that network parameters, like the total number of nodes M, are known. This assumption might be valid in some networks operated by a single authority (e.g., sensor networks). Nevertheless, in many envisioned applications such parameters might be unknown (e.g., a vehicle that just connected to a local VANET). In order to make Spray and Wait equally efficient in such scenarios as well, we would like to produce and maintain good estimates of necessary network parameters, like M, and adapt L accordingly.

This problem is difficult in general. A straightforward way to estimate M would be to count unique IDs of nodes encountered already. However, this method requires a large database of node IDs to be maintained in large networks, and a lookup operation to be performed every time any node is encountered. Furthermore, although this method converges eventually, its speed depends on network size and could take a very long time in large disconnected networks. A better alternative is to produce an estimate of M by taking advantage of inter-meeting time statistics. Specifically, let us define T_1 as the time until a node (starting from the stationary distribution) encounters *any* other node. It is easy to see from Lemma 4.2 that T_1 is exponentially distributed with average $T_1 = EM_{mm}/(M-1)$. Furthermore, if we similarly define T_2 as the time until two *different* nodes are encountered, then the expected value of T_2 equals $EM_{mm}(\frac{1}{M-1} + \frac{1}{M-2})$. Cancelling EM_{mm} from these two equations we get the following estimate for M:

$$\hat{M} = \frac{2T_2 - 3T_1}{T_2 - 2T_1}. \quad (3)$$

Estimating M by the procedure above presents some challenges in practice, because T_1 and T_2 are ensemble averages. Since hitting times are ergodic [41], a node could collect sample intermeeting times $T_{1,k}$ and $T_{2,k}$ and calculate time averages $\hat{T}_1$ and $\hat{T}_2$ instead. However, when a node i meets another node j, i and j become *coupled* [43]; in other words, the next intermeeting time of i and j is not anymore exponentially distributed with average EM_{mm}. In order to overcome this problem, each node keeps a record of recently encountered nodes. Every time a new node is encountered, it is stamped as "coupled" for an amount of time equal to the *mixing* or *relaxation* time for that graph, which is the expected time until a node starting from a given position arrives to its stationary distribution [41].[4] Then, when node i measures the next sample intermeeting time, it ignores all nodes that it is coupled with at the moment, denoted as c_k, and scales the collected sample $T_{1,k}$ by $\frac{M-c_k}{M-1}$. A similar procedure is followed for $\hat{T}_2$. Putting it altogether, after n samples have been collected:

$$\hat{T}_1 = \frac{1}{n}\sum_{k=1}^{n}\left(\frac{M-c_k}{M-1}\right)T_{1,k},$$

$$\hat{T}_2 = \frac{1}{n}\sum_{k=1}^{n}\left[\left(\frac{M-c_{k-1}}{M-1}\right)T_{1,k-1} + \left(\frac{M-c_k}{M-2}\right)T_{1,k}\right].$$

Replacing $\hat{T}_1$ and $\hat{T}_2$ in (3) we get a current estimate of M. As can be seen by (3), the estimator for M is sensitive to small deviations of T_1 and T_2 from their actual values. Therefore, it is useful for a node to also maintain a running average of M. Specifically, the running

[3]By this, we do not assume that ED_{opt} is always known to the user. If ED_{opt} is not known a could still be used as a measure of how "aggressive" the protocol should be.

[4]If nothing about the network size and mobility model is known either, a node could alternatively set this parameter arbitrarily to a large value, and possibly adapt it according to frequency by which a node with the same ID is encountered.

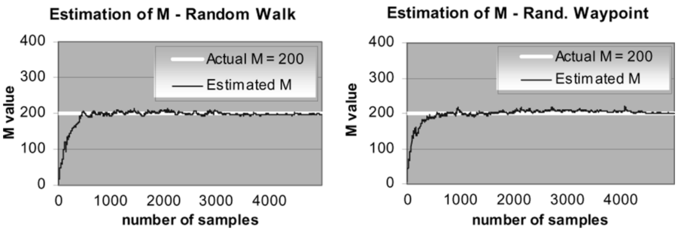

Fig. 3. Online estimator of number of nodes (M): $N = 200 \times 200$, transmission range $= 0$, $\beta = 0.98$, mixing time $= 4000$.

estimate $\hat{M}$ is updated with every new estimate $\hat{M}_{new}$ as $\hat{M} = \beta \hat{M} + (1-\beta)\hat{M}_{new}$ $(0 < \beta < 1$, with values closer to 1 providing better stability). We could now use this estimate of M to calculate the number of copies using one of the previous methods.

Fig. 3 shows how the online estimate $\hat{M}$, calculated with our proposed method, quickly converges to its actual value for a 200 $\times$ 200 network with 200 nodes, for both the random walk and random waypoint models, again validating the generality of our expressions. (Note that even in this small scenario, our method's convergence is more than two times faster than ID-counting.) We believe that similar estimators could potentially be constructed for other network parameters or statistics, as well, (e.g., higher moments for encounter times) which could be used to provide users with predictions of the service level available. Finally, both our method and ID-counting could take advantage of indirect information learning, where nodes exchange known unique IDs or independently collected samples to speed up convergence.

C. Scalability of Spray and Wait

Having shown how to find the minimum number of copies L_{min} to achieve a delay at most a times the optimal, it would be interesting, from a scalability point of view, to see how the percentage L_{min}/M of nodes that need to receive a copy behaves as a function of M. The reason for this is the following: If we assume a large enough TTL (time-to-live) value, flooding-based schemes will eventually give a copy to every node (i.e., at least M transmissions). Increased contention and the resulting retransmissions increase this value significantly, as we shall see. On the other hand, Spray and Wait performs L transmissions, and produces very little contention. Consequently, the number of transmissions that Spray and Wait performs per message is at most a fraction L_{min}/M of the number of transmissions per message that epidemic and other flooding-based schemes perform.

Lemma 5.2: Let L/M be constant and let $L \ll M$. Let further $L_{min}(M)$ denote the minimum number of copies needed by Spray and Wait to achieve an expected delay that is at most aED_{opt}, for some a. Then $\frac{L_{min}(M)}{M}$ is a decreasing function of M.

Proof: When $L \ll M$ we can use the upper bound of (2) to examine the behavior of Spray and Wait:

$$ED_{sw} \leq EM_{mm}(H_{M-1} - H_{M-L}) + \left(\frac{M-L}{M-1}\right)\frac{EM_{mm}}{L}.$$

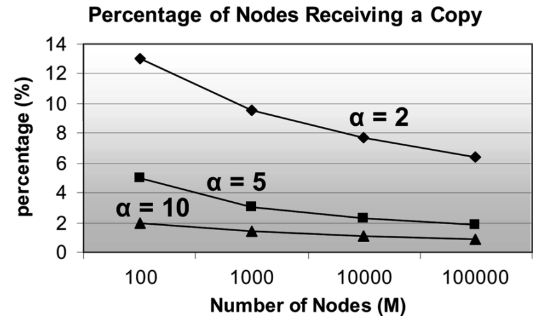

Fig. 4. Required percentage of nodes L_{min}/M receiving a copy for spray and wait to achieve an expected delay of aED_{opt}.

Since $H_n = \Theta(\log(n))$, $H_{M-1} - H_{M-L} = \Theta(\log(\frac{M-1}{M-L}))$. Also, let $L = \gamma M$, where γ is a constant $(\gamma \ll 1)$. Replacing L in the previous equation gives us

$$\Theta\left(\log(1-1/M) - \log(1-\gamma)\right) EM_{mm} + \left(\frac{M}{M-1}\right)\left(\frac{1-\gamma}{\gamma}\right)\frac{EM_{mm}}{M}.$$

Now, for large M, $\frac{M-1}{M} \simeq 1$. Therefore, keeping the size of the grid N and transmission range K constant, we get that $ED_{sw} = \Theta(1) + \Theta(1)(\frac{1}{M}) = \Theta(\frac{1}{M})$.

On the other hand, for constant N and K, $ED_{opt} = \Theta(\frac{\log(M)}{M})$ as was shown in [1]. Hence, $ED_{sw}/ED_{opt} = \Theta(\frac{1}{\log(M)})$ (i.e., decreasing with M), if L/M is kept constant. This implies that if we require ED_{sw}/ED_{opt} to be kept constant for increasing M, then L/M has to be decreasing. ■

What this interesting result says is the following: *If we keep the relative transmission overhead (i.e., transmissions/msg) between Spray and Wait and the Optimal scheme constant, then the relative delay of Spray and Wait, again compared against the optimal, improves as the number of nodes increase.* Alternately, if we keep the relative delay constant, then Spray and Wait would require a smaller and smaller percentage of the total nodes to act as relays for a given message. In other words, Spray and Wait benefits from a higher number of nodes more than the Optimal scheme does. This behavior implies that Spray and Wait is extremely scalable, unlike flooding-based schemes. In Fig. 4 we depict the behavior of L_{min}/M as a function of M for different values of a.

Remark: In presenting the methods and algorithms of this section we have focused on Spray and Wait. The reason for this is that we do not possess an analytical expression for the delay of Spray and Focus. However, since the delay of Spray and Wait is an upper bound on the delay of Spray and Focus for many mobility models, the same algorithms could be used also to provide a pessimistic estimate (upper bound) on the number of copies that Spray and Focus needs to achieve a given delay. In the next section, we show that a better rule-of-thumb for Spray and Focus is to use 1/4 to 1/2 of the copies calculated for Spray and Wait.

VI. Simulation Results

We have used a custom discrete event-driven simulator to evaluate and compare the performance of different routing protocols under a variety of mobility models and under contention.

MAC Protocol and Channel Model: A very simplified version of a slotted CSMA (Carrier-Sense Multiple Access) MAC protocol has been implemented. Each message takes one time slot to be transmitted, and the channel is sensed at the beginning of the slot. Thus, contention is avoided at the sender, but collisions may occur at the receiver ("hidden terminal"). If a message is received successfully, a small (link-layer) acknowledgement packet (ACK) is sent back to the sender. Also, if more than one messages need to be sent, they are sent (and ACKed) in a burst. Although we expect that the choice of MAC protocol will have some impact on performance, it is beyond the scope of this paper to cover all possible options. Further, we believe that a simple CSMA MAC is a fairly realistic option in this context.[5] Finally, we adopt a simple channel model, where transmission coverage is circular, and interference may occur only from immediate neighbors. Yet, the routing protocols do not make any assumptions about the physical layer, and more detailed channel models could be used. If the channel conditions are good, a link is established, otherwise not.

The Routing Protocols: We have implemented and compared the following routing protocols (we will use the shorter names in the parentheses to refer to each routing scheme in simulation plots):

Epidemic routing ("epidemic"): a node copies a message to every new node it encounters that does not have a copy already. For this and all other protocols, we choose TTL (time-to-live) values between 1000–10000 time units for each message.

Randomized flooding or Gossiping ("random-flood"): like epidemic routing, but a message only gets copied with some probability $p < 1$ (we have used values between 0.5 and 0.05).

Utility-based flooding ("utility-flood"): like epidemic routing, but a message gets copied only if the node encountered has a utility value higher than the current by some threshold U_{th} (we have used the utility function from [17] and the values for U_{th} are between 10 and 90).

Binary Spray and Wait ("spray&wait"): We choose the number of copies L to be equal to about 10%–15% of all nodes M, according to the theory of Section V,

Spray and Focus ("spray&focus"): We have found that choosing L equal to about 5%–10% of the total nodes serves as a useful rule of thumb for good performance.

Seek and Focus single-copy routing ("seek&focus") [1]: We have also included the champion scheme from our single-copy study in one scenario. Seek and Focus forwards the (single) copy randomly at the beginning, until a node with a high enough utility is found. Then, it switches to utility-based forwarding to route the (single) copy towards its destination. More details about its mechanism can be found in [1].

We first evaluate the effect of traffic load, bandwidth, and storage capacity on the performance of different routing schemes. We then examine their performance as the level of connectivity or mobility model changes.

[5]Collision avoidance (CSMA/CA), for example, may be undesirable due to the extra overhead of RTS,CTS packets. Nevertheless, to ensure that any conclusions drawn are not mere artifacts of our choice of MAC, we have also run some simulations with an ideal CSMA/CA protocol, as well as for a "broadcast" CSMA MAC (i.e., without acknowledgements). Although the numbers do differ, we observed that the relative performance between the protocols is largely unchanged.

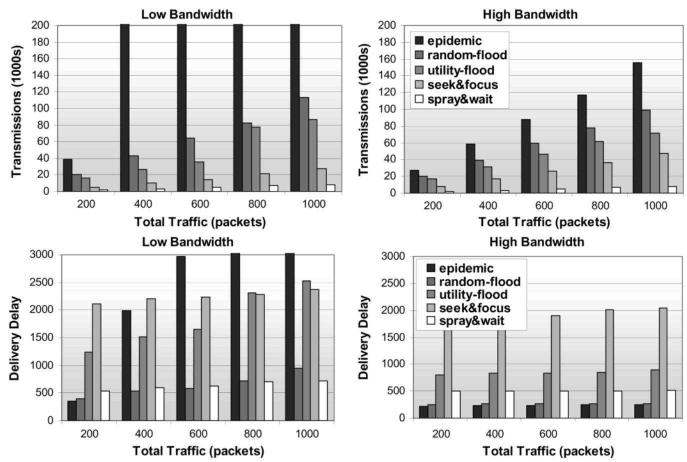

Fig. 5. Total transmissions (top) and delivery delay (bottom) of all routing protocols under varying traffic loads. The plots on the left correspond to low link bandwidth, and the plots on the right to high bandwidth.

A. *Effect of Traffic Load, Bandwidth, and Storage Capacity*

One hundred nodes move according to the random waypoint model [44] in a 500 × 500 grid with reflective barriers. The transmission range K of each node is equal to 10. Each node selects a destination randomly and starts generating messages for it until time 10 000 (the simulation ends at time "$10\,000 + TTL$" to give an equal chance to messages generated towards the end to be delivered). Finally, the message inter-arrival time is uniformly distributed in $[0, T_{max}]$ with T_{max} from 10 000 (low traffic—around 200 messages in total) to 2000 (high traffic—around 1000 messages in total).

Fig. 5 depicts the total number of transmissions and average delivery delay for all routing algorithms. Additionally, it does so for two different values of link bandwidth assumed. In the first case (*low bandwidth*) the bandwidth of a given contact might not be enough to forward all intended messages. (Note that time and link capacity is normalized with one message transmission taking one time unit.) In the second case (*high bandwidth*), the link bandwidth is 4 times larger, and does not become the bottleneck for the traffic loads considered. (We do not include results for Spray and Focus in this scenario as it had very similar performance with Spray and Wait.)

As is evident by Fig. 5, Spray and Wait performs significantly fewer transmissions (up to $10\times$) than all single and multi-copy protocols, under all conditions. In terms of delivery delay, if traffic loads are low or network bandwidth ample (epidemic has close-to-optimal delays under these conditions), it manages to achieve delays that are quite close to those of flooding-based schemes . What is more, if traffic starts increasing or the available bandwidth is reduced, it actually outperforms all schemes in terms of delay also (up to $3\times$ improvement). (Note that almost all schemes in this scenario had delivery ratios above 90%, except Seek and Focus which had about 70%, and Epidemic routing which plummeted to less than 50% for very high traffic in the low bandwidth scenario.). The above results imply that bandwidth has an important effect on the performance of different protocols. Even though the actual values, like contact

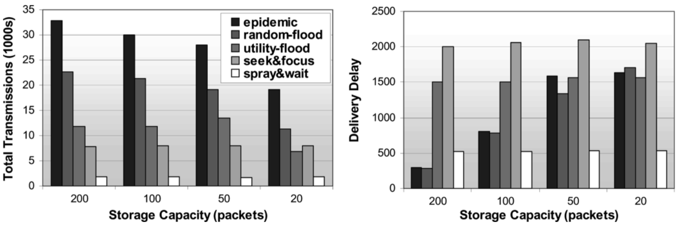

Fig. 6. Performance comparison of all routing protocols as a function of node storage capacity.

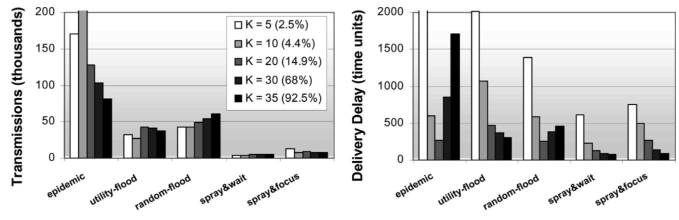

Fig. 7. Random Waypoint Mobility: total transmissions and delay as a function of transmission range K (respective connectivity values are shown in parentheses).

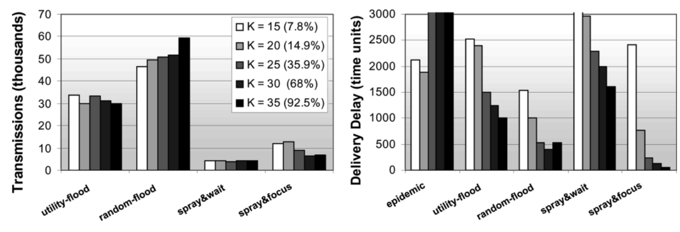

Fig. 8. Random Walk Mobility: total transmissions and delay as a function of transmission range K (respective connectivity values are shown in parentheses).

time, link rate, message size, etc. depend on a number of application-specific parameters, we can safely conclude the following: if the available network bandwidth is much higher than the total traffic load to be accommodated, then flooding-based schemes are quite fast but spraying schemes can deliver comparable delays with much fewer transmissions; if bandwidth becomes limited then flooding-based schemes suffer from contention and their delay is also higher than that of Spray routing. A more detailed analysis of the effect of bandwidth on the performance of epidemic routing can be found in [38].

A similar effect is in work also with buffer capacity. When a new node is encountered, who has little remaining buffer space, not all messages that could otherwise be forwarded actually do. This, as in the case of lack of bandwidth, results in extra queueing delay for the unlucky messages (which are more numerous in the case of flooding schemes). Fig. 6 compares the total transmissions and delays of all protocols, as the node storage capacity decreases from 200 message copies to 20 (the traffic load assumed is 200 messages in total, generated according to the previously described method). As can be seen there, Epidemic routing and Random flooding start to suffer from increasing queueing delays as buffer size shrinks, while Spray and Wait (as well as Utility Flooding and Seek and Focus) sees no significant performance change. (Note that we assume that a message can be forwarded to its destination, even if its buffer is full, in order to allow buffers to eventually drain; thus, delivery ratios were high for all protocols). For a more detailed evaluation of the effect of limited buffer space on the performance of flooding one can look in [12], [13], and [15], where similar conclusions can be drawn.

B. *Effect of Connectivity and Mobility Model*

In this scenario, we would like to evaluate the performance of all protocols in networks with a large range of connectivity characteristics, ranging from very sparse, highly disconnected networks, to *almost* connected networks. We assume a medium traffic load (500 messages), generated as described in the previous scenario.

Before we proceed, we need to define a meaningful connectivity metric. Although a number of different metrics have been proposed (for example [45]), 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 ("giant component"). This indicates what percentage of nodes have already conglomerated into the connected part of the network, with "one" implying a regular connected network (with high probability). We vary the transmission range K to span the entire connectivity range.

The above connectivity metric measures "static" connectivity. It indicates how connected a random snapshot of the connectivity graph will be. However, in situations where mobility is exploited to deliver traffic end-to-end, "dynamic" connectivity also plays an important role on performance. Dynamic connectivity can be seen as a measure of how many new nodes are encountered by a given node within some time interval. If nodes move in an IID manner, it is directly tied to the mixing time for the graph representing the network [41]. 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.

In order to evaluate the effect of dynamic connectivity on different protocols, we present three sets of results, one where nodes move according to the Random Waypoint model, one where nodes perform Random Walks, and one where nodes move according to a "Community-based" mobility model that tries to capture some important mobility characteristics observed in real traces.

Let us first look at the two popular mobility models, random waypoint and random walk mobility. The random waypoint has one of the fastest mixing times $(\Theta(\sqrt{N}))$, while the random walk has one of the slowest $(\Theta(N))$ [41]. In both cases, we assume there are 100 nodes in a 200×200 network.

Fig. 7 and Fig. 8 depict the number of transmissions and the average delay for the random waypoint and the random walk scenarios, respectively, as a function of transmission range (respective connectivity values are shown in the parentheses). We only depict results for multi-copy schemes here (a detailed treatment of single-copy routing performance as a function of connectivity can be found in [17]).

There are a number of interesting things to notice about these plots. First, although Randomized and Utility Flooding can improve the performance of epidemic routing they still have to perform way too many transmissions to achieve competitive delays. Further, when nodes move according to the random waypoint model, Spray and Wait outperforms all protocols, in terms of both transmissions and delay, for all levels of connectivity. Its performance is close to the optimal, and thus Spray and Focus cannot offer any improvement.

On the other hand, when nodes perform random walks, Spray and Wait may exhibit large delays, if the network area is large. Here the few copies are spread locally, and then each relay takes a long time to traverse the network and reach the destination. Even if the number of copies were increased, it would be the spraying phase that would take a long time, since new nodes are found very slowly. (Note though that the delivery ratio for Spray and Wait did not ever fall under 90%). Spray and Focus can overcome these shortcomings and excel (when the network is not too sparse), achieving the smallest delay with only a few extra transmissions. Note also that, despite using the same utility function as Spray and Focus, Utility Flooding is still plagued by its flooding nature. This problem was even more pronounced when other existing utility functions were used [14].

Finally, epidemic routing and the rest of the schemes manage to achieve good delays for a few connectivity values, but perform poorly (and nonlinearly) for most values. Spray and Wait and Spray and Focus, on the other hand, exhibit greater stability. They perform few transmissions across all scenarios, while achieving a delivery delay that decreases as the level of connectivity increases, as one would expect.

Community-Based Mobility: Popular mobility models like the ones we have 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 from all others. 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 [32], [33]. 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 [42].

In the Community-based model, each node has its own small community ($c\times$ the size of the network, $c < 1$) inside which it moves preferentially for the majority of time (e.g., the user's department building on a campus). Every now and then it leaves its community and roams around the network for sometime (e.g., going to a class at a different building, to a dining hall, library, etc.), and then returns. Finally, each node may have different mobility characteristics in addition to different communities. Some nodes may spent a very large amount of their time inside their community, while others may be more "mobile". This Community-based model allows for a large range of node heterogeneity to be captured.

In Fig. 9, we depict the total transmissions and average delivery delay for the Community-based model ($c = 1/25$) with heterogeneous node mobility. Specifically, a node leaves its community after a "local trip" with a probability p_l chosen uniformly in [0.05, 0.4]. p_l is chosen at the beginning of the scenario independently for each node. Further, the probability p_r that a node returns to its community after a "roaming trip" is chosen uniformly in [0.8, 0.4] for each node.[6] The conclusions that can be drawn from these plots are similar to those in the previous two scenarios. Specifically, because there are some nodes in the network that are quite mobile (unlike the Random Walk case), Spray and Wait manages to achieve good performance. (However, note that in this case not all relays are equally useful to the delivery process; only the more mobile relays are the ones that are mostly successful in delivering messages for Spray and Wait.) Yet, Spray and Focus can again take advantage of the high locality of many nodes, and deliver messages 5–6 times faster, especially as density increases (i.e., when transmission range increases).

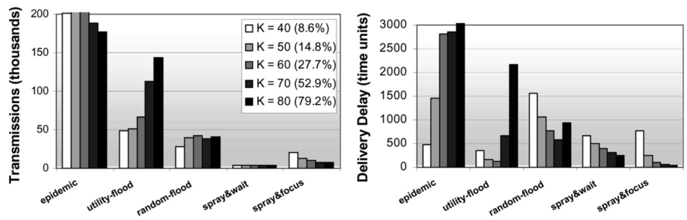

Fig. 9. Community-based Mobility: Total transmissions and delay as a function of transmission range K (respective connectivity values are shown in parentheses).

VII. Conclusion

In this work, we investigated the problem of multi-copy routing in intermittently connected mobile networks. We proposed two efficient multi-copy schemes, called Spray and Wait and Spray and Focus, that manage to overcome the shortcomings of flooding-based and other existing schemes. Using theory and simulations we showed that: (i) when enough nodes in the network are sufficiently mobile, Spray and Wait outperforms existing schemes with respect to both number of transmissions and delivery delays, and achieves comparable delays to an optimal scheme, despite its simplicity, and (ii) when node mobility is low or predominantly local, Spray and Focus can retain the performance advantage of Spray and Wait with only a small overhead on total transmissions and simplicity. Finally, both schemes are very robust to network size and density changes.

In future work, we intend to extend our analysis to cover contention for the wireless channel, and more realistic mobility models that might exhibit correlation in space and time.

References

[1] 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.

[6]We assume that mobility here consists of trips or "epochs", as in the random waypoint or random direction case, with the difference that now a two-state Markov Chain decides whether the next trip will be a local one (inside the community) or a roaming one.

[2] P. Gupta and P. Kumar, "Capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, 2000.
[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] 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.
[5] Disruption Tolerant Networking. [Online]. Available: http://www.darpa.mil/ato/solicit/DTN/
[6] 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.
[7] J. Heidemann, W. Ye, J. Wills, A. Syed, and Y. Li, "Research challenges and applications for underwater sensor networking," in *Proc. IEEE Wireless Communications and Networking Conf.*, 2006.
[8] Delay Tolerant Networking Research Group. 2007 [Online]. Available: http://www.dtnrg.org
[9] J. Scott, P. Hui, J. Crowcroft, and C. Diot, "Haggle: A networking architecture designed around mobile users," in *Proc. IFIP Conf. Wireless On-Demand Network Systems and Services (WONS)*, 2006.
[10] M. Papadopouli and H. Schulzrinne, "Seven degrees of separation in mobile ad hoc networks," in *Proc. IEEE GLOBECOM*, 2000.
[11] C. E. Perkins, *Ad Hoc Networking*, 1st ed. Reading, MA: Addison-Wesley, 2001.
[12] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Duke Univ., Durham, NC, Tech. Rep. CS-200006, Apr. 2000.
[13] 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.
[14] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic routing in intermittently connected networks," *SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 7, no. 3, 2003.
[15] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," in *Proc. IFIP Networking*, 2006.
[16] J. Widmer and J.-Y. Le Boudec, "Network coding for efficient communication in extreme networks," in *Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (WDTN)*, 2005.
[17] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: The single-copy case," *IEEE Trans. Networking*, vol. 16, no. 1, Feb. 2008.
[18] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *Proc. MobiHoc*, 2004.
[19] 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.
[20] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Proc. ACM SIGCOMM*, Aug. 2004.
[21] 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.
[22] 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.
[23] 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.
[24] 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.
[25] 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.
[26] 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.
[27] 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.
[28] F. Tchakountio and R. Ramanathan, "Tracking highly mobile endpoints," in *Proc. ACM Workshop on Wireless Mobile Multimedia (WOWMOM)*, 2001.
[29] G. Sharma and R. Mazumdar, "On achievable delay/capacity trade-offs in mobile ad hoc networks," in *Proc. IEEE WiOpt*, 2004.
[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. Jain, M. Demmer, R. Patra, and K. Fall, "Using redundancy to cope with failures in a delay tolerant network," in *Proc. ACM SIGCOMM*, 2005.
[32] 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.
[33] 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.
[34] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft, "Opportunistic content distribution in an urban setting," in *Proc. ACM CHANTS*, 2006.
[35] 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.
[36] 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.
[37] 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.
[38] A. Jindal and K. Psounis, "Performance analysis of epidemic routing under contention," in *Proc. ACM IWCMC*, 2006.
[39] R. Groenevelt, G. Koole, and P. Nain, "Message delay in Manet (extended abstract)," in *Proc. ACM Sigmetrics*, 2005.
[40] 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.
[41] 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
[42] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Performance analysis of mobility-assisted routing," in *Proc. ACM/IEEE MOBIHOC*, 2006.
[43] R. Durrett, *Probability: Theory and Examples*, 2nd ed. London, U.K.: Duxbury Press, 1995.
[44] 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.
[45] O. Dousse, P. Thiran, and M. Hasler, "Connectivity in ad hoc and hybrid networks," in *Proc. IEEE INFOCOM*, 2002, pp. 1079–1088.

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.