

Optimizing Multi-Copy Routing Schemes for Resource Constrained Intermittently Connected Mobile Networks

Apoorva Jindal

Department of Electrical Engineering
University of Southern California
Los Angeles, California 90089
Email: apoorvaj@usc.edu

Konstantinos Psounis

Department of Electrical Engineering
University of Southern California
Los Angeles, California 90089
Email: kpsounis@usc.edu

Abstract— Intermittently connected mobile networks are wireless networks where most of the time there does not exist a complete path from a source to a destination. Researchers have proposed flooding based schemes for routing in such networks. While flooding based schemes are robust and have a high probability of delivery, they suffer from a huge overhead in terms of bandwidth, buffer space and energy dissipation due to large number of transmissions per packet. So flooding based schemes are impractical for resource constrained networks. Controlled replication or spraying methods can reduce this overhead by distributing a small, fixed number of copies to only a few relays, which then independently route each copy towards the destination. These schemes demonstrate a good delay performance without using a lot of resources.

There are three important questions in the context of the design of these spraying based schemes: (i) How many copies per packet should be distributed? (ii) How to distribute these copies amongst the potential relays? (iii) How are each of these copies routed towards the destination? The first and the third questions have been studied in detail by different researchers. But, there has been no study which looks at the second question. This paper fills this void. Specifically, we propose a methodology to derive the optimal spraying policy. As a case study, we find the optimal spraying policy for two different spraying based schemes. Finally, we study the optimal policies to infer simple heuristics which achieve expected delays very close to the optimal.

I. INTRODUCTION

Intermittently connected mobile networks (ICMNs) are networks where most of the time there does not exist a complete end-to-end path from a source to a destination. Even if such a path exists, it is highly unstable and may change or break soon after it has been discovered. This situation arises when the network is quite sparse. Examples of such networks include wildlife tracking and habitat monitoring sensor networks [1], low cost Internet provisioning to remote communities [2], [3], military networks [4], inter planetary networks [5], vehicular ad hoc networks [6] etc.

Conventional mobile ad hoc routing protocols, such as DSR, AODV etc. [7], assume that a complete path exists between a source and a destination, and try to find minimum cost paths before any useful data is sent. Since no such end-to-end paths exist most of the time in ICMN's, conventional protocols will

fail to deliver any data to all but the few connected nodes. To overcome the disconnectedness of connectivity graph, a new model of routing has been proposed. Mobility of nodes is exploited to route packets towards the destination. A node stores the packet for long periods of time as it moves around before an appropriate communication opportunity arises. Routing here consists of independent, local forwarding decisions, based on current connectivity information and predictions of future connectivity information. The crucial question any routing algorithm has to answer in this context is 'what is a good next hop when no path to the destination is currently available'?

The answer to this question by the majority of existing protocols has been 'everyone' (flooding or epidemic routing [8]) or 'almost everyone' (randomized flooding and utility based flooding [9], [10]). The main problem with flooding based schemes is the overhead involved in terms of bandwidth, buffer space and energy dissipation due to the large number of transmissions per packet. So flooding based schemes are inappropriate for resource constrained networks. To solve this problem, researchers have proposed the use of controlled replication or *spraying* [9], [11]–[14]. A fixed, small number of copies are generated and distributed to different relays, each of which is then independently routed towards the destination. By routing multiple copies independently, these schemes create enough diversity to explore the sparse connectivity graph efficiently, while keeping the resource usage per packet low.

The three important questions in the context of the design of spraying based routing schemes are: (i) How many copies per packet should be distributed? The answer to this question depends on the resource constraints in the network and has been studied by [11]. (ii) How to distribute these copies amongst the potential relays? This is still an open question and is the focus of this paper. (iii) Once the copies have been distributed, how should each of these copies be routed towards the destination? Two different schemes to route the individual copies have been proposed in the literature: (i) Spray and Wait [11]: Each relay node uses direct transmission to route the packet to the destination, that is, the relay forwards the copy to

the destination only. (ii) Spray and Focus [13]: Each relay node performs utility based forwarding towards the destination, that is, if the relay node meets another node having a higher utility, it gives its copy to the other node. Spray and Focus has a better delay than Spray and Wait but performs a lot more transmissions and hence consumes more energy.

This paper studies how to distribute the copies amongst the potential relays such that the expected end-to-end delay is minimized. Previous spraying based schemes [11], [13] use binary spraying to distribute the copies: If a node which has more than one copy encounters a potential relay, it hands over half of its copies to the relay. If all the relays are identical, then choosing one over the other does not give any advantage, and binary spraying will have the best performance amongst all spraying schemes [13] because it sprays copies whenever it gets an opportunity to do so. But, if the utilities of each node are known, then relays with higher utilities should be more preferable. In such a scenario, if a node encounters a lower utility node, it has to decide whether to give it some of its copies or save its copies for a node with higher utility which it may encounter in future.

This paper presents a methodology to find the optimal spraying scheme. The methodology involves solving a system of non linear equations. We use approximations to simplify these equations, which can then be solved using a dynamic program. Then as case studies, we find the optimal spraying scheme for both Spray and Wait and Spray and Focus. We study the optimal spraying policy to infer simple heuristics which achieve expected delays very close to the optimal.

The outline of the paper is as follows: First Section II presents our notation and assumptions, and then states some useful results for the random walk mobility model which we use during the remainder of the paper. Then, Section III presents the methodology to find the optimal spraying scheme. Section IV uses this methodology to find the optimal spraying policies for Spray and Wait and Spray and Focus. Finally, Section V concludes the paper.

II. PRELIMINARIES

A. Notation and Assumptions

We first introduce our notation and state the assumptions we will be making throughout the remainder of the paper.

- (i) M nodes perform independent random walks on a $\sqrt{N} \times \sqrt{N}$ 2D torus (finite lattice). Each node moves one grid unit in one time unit.
- (ii) Each node can transmit up to $K \geq 0$ grid units away, where $\frac{K}{\sqrt{N}}$ is much smaller than the value required for connectivity [15]. We use Manhattan distance $d_{ab} = \|a_x - b_x\| + \|b_x - b_y\|$ to measure proximity between two positions a and b (or between two nodes).
- (iii) There is no contention in the network. In other words, the buffer space is infinite, and any communicating pair of nodes do not interfere with any other simultaneous transmission.
- (iv) Let the number of copies distributed by the spraying based schemes be denoted by L .

B. Useful Results for Random Walk Mobility Model

This section presents some results on random walks which will be used during the remainder of the paper.

Lemma 2.1: Let $E[M]$ denote the expected time until two independent random walks, starting from their stationary distribution, first meet each other. It is also referred to as the expected meeting time and is equal to $\frac{N}{2} \left(\log(N) - \frac{2^{K+1}-K-2}{2^K-1} \right)$.

Proof: See [16] \square

Lemma 2.2: Let $E[M(d)]$ denote the expected time until two independent random walks, starting at a distance d from each other, first meet each other. $E[M(d)]$ can be derived by solving the following set of linear equations:

$$E[M(d)] = \begin{cases} p_{d,d-2}E[M(d-2)] + p_{d,d}E[M(d)] + p_{d,d+2}E[M(d+2)] & d > K \\ 0 & d \leq K \end{cases}, \quad (1)$$

where p_{d_1,d_2} denotes the probability that the two walks are at a distance d_2 from each other in the next time slot given they are at a distance d_1 from each other in the current time

slot and, for $d_1 > 3$, it equals $\begin{cases} \frac{16d_1-20}{64d_1} & d_2 = d_1 - 2 \\ \frac{16d_1+12}{64d_1} & d_2 = d_1 + 2 \\ \frac{32d_1+8}{64d_1} & d_2 = d_1 \end{cases}$,

for $d_1 = 3$, it equals $\begin{cases} \frac{7}{48} & d_2 = 1 \\ \frac{15}{48} & d_2 = 5 \\ \frac{28}{48} & d_2 = 3 \end{cases}$,

for $d_1 = 2$, it equals $\begin{cases} \frac{3}{32} & d_2 = 0 \\ \frac{11}{32} & d_2 = 4 \\ \frac{18}{32} & d_2 = 2 \end{cases}$, for $d_1 = 1$, it equals $\frac{7}{16}$ and $\frac{9}{16}$ for $d_2 = 3$

and $d_2 = 1$ respectively and for $d_1 = 0$, it equals $\frac{4}{16}$ and $\frac{12}{16}$ for $d_2 = 2$ and $d_2 = 0$ respectively.

Sketch of Proof: If $d \leq K$, then $E[M(d)] = 0$ trivially. If $d > K$, then in the next time slot, the two nodes can either move closer by two grid units or move further away by two grid units or still remain at the same distance. (Since each node can move only one grid unit in one time unit.) Thus, Equation (1) follows directly from the law of total probability. The value of p_{d,d_2} can be evaluated using elementary combinatorics. \square

III. OPTIMAL SPRAYING SCHEME

This section presents the methodology to find the optimal spraying policy. Specifically, the algorithm will answer the following question: ‘Two nodes A and B are within range of each other and A has l copies of a packet while B has none. The utility of both the nodes is known. Then how many of the l copies should A give to B such that the expected delivery delay is minimized.’

Before we proceed, we first specify the utility function we will use. Amongst the different utility functions used in the literature (see [16]), we choose ‘the distance to the destination’ for our analysis.

Now we derive the algorithm to find the optimal spraying policy. Let a node (label it node A) be at a distance d from the destination and has l copies of the packet. Let $D(d, l)$ denote the time this node will take to deliver the packet to

the destination. In the future time slots, either one of the following two events can happen first: (i) E_1 : Node A meets the destination and delivers the packet. (ii) E_2 : Node A meets one of the potential relays. Let the time duration elapsed till event E_i occurs be denoted by T_i , $i = 1, 2$. By definition, T_1 is exponentially distributed with mean $E[M(d)]$. To derive the distribution of T_2 , we use the fact that the time it takes to meet one particular relay node is exponentially distributed with mean $E[M]$. T_2 is the minimum of M^1 such exponentials which is also an exponential with mean $\frac{E[M]}{M}$. Thus, the time duration till one of these two events occur is equal $\min(T_1, T_2)$ and is exponentially distributed with mean $\frac{1}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}}$.

Let node A encounter a potential relay (lets label it node B) before meeting the destination. (The probability of this event is equal to $\frac{\frac{M}{E[M]}}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}}$.) Let node A and B be at a distance d_A and d_B from the destination when they meet. Node A has l copies of the packet while B has none. Let $D_M(d_A, d_B, l)$ denote the minimum additional delay to deliver the packet to the destination. Then,

$$E[D(d, l)] = \frac{1}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} + \frac{\frac{M}{E[M]}}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} \sum_{d_A, d_B} P(d_A, d_B) E[D_M(d_A, d_B, l)], \quad (2)$$

where $P(d_A, d_B)$ is the probability that the two nodes are at a distance d_A and d_B from the destination when they meet.

Node A can give any number from 0 to $l - 1$ copies to the B . If i of the l copies are given to B , then the delivery delay to the destination is the minimum of $D(d_A, l - i)$ and $D(d_B, i)$. Hence,

$$E[D_M(d_A, d_B, l)] = \min_{0 \leq i < l} (E[\min(D(d_A, l - i), D(d_B, i))]) \quad (3)$$

Note that the solution to Equation (3) gives the optimal spraying policy.

Equations (2) and (3) form a system of non linear equations. Solving these equations will give the optimal spraying policy, but solving a non linear system is not easy. So, we make approximations to simplify these equations. (Note that due to these approximations, the spraying policy obtained is not really the optimal, but it will give an intuition into the structure of the optimal policy.)

First, we assume that the sum of two exponentially distributed random variables is also exponential. With this approximation, the distribution of both $D(d, l)$ and $D_M(d_A, d_B, l)$ can be derived to be exponential. Thus, Equa-

tion (2) reduces to the following:

$$E[D(d, l)] = \frac{1}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} + \frac{\frac{M}{E[M]}}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} \sum_{d_A, d_B} P(d_A, d_B) \min_{0 \leq i < l} \left(\frac{1}{\frac{1}{E[D(d_A, l-i)]} + \frac{1}{E[D(d_B, i)]}} \right). \quad (4)$$

Equation (4) is still a system of non linear equations which are not easy to solve. So, we make another approximation by replacing d_A by its expected value. For the random walk mobility model, $E[d_A]$ is equal to d as the probability of moving in any direction is the same. Replacing d_A by d in Equation (4) yields,

$$E[D(d, l)] = \frac{1}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} + \frac{\frac{M}{E[M]}}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} \sum_{d_B} P(d_B | d_A = d) \min_{0 \leq i < l} \left(\frac{1}{\frac{1}{E[D(d, l-i)]} + \frac{1}{E[D(d_B, i)]}} \right). \quad (5)$$

In Equation (5), the value of $E[D(d, l)]$ depends only on those $E[D(d, l)]$ for which either $l < l$ or $l = l, d \leq d$. So, a dynamic program can be used to solve Equation (5).

The dynamic program will be initialized with the value of $E[D(d, 1)]$ which depends on how each copy is routed towards the destination. Section IV finds its value for Spray and Wait and Spray and Focus.

The only unknown in Equation (5) is $P(d_B | d_A = d)$. Since node B is within range of A , d_B will lie within $d - K$ and $d + K$. $P(d_B | d_A = d)$ can be derived using elementary combinatorics to be equal to

$$\begin{cases} \frac{K+1}{4K} & d_B = d - K \\ \frac{4K}{2} & d - K + 2 \leq d_B \leq d + K - 2 \\ \frac{K+1}{4K} & d_B = d + K \end{cases}.$$

IV. CASE STUDIES

This section studies the optimal spraying policy for two different copy routing strategies. For each strategy, we first find $E[D(d, 1)]$, then study the spraying policy obtained by solving Equation (5) and finally present a simple heuristic which achieves a expected delay very close to the optimal.

A. Spray and Wait

In Spray and Wait, each relay node routes the copy towards the destination using direct transmission. Thus, $E[D(d, 1)]$ is the expected time it takes for the relay to meet the destination and is equal to $E[M(d)]$.

Now, we study the spraying policy obtained by solving Equation (5). Let node A which has l copies of the packet meet node B which has none. Let the distance to the destination of both the nodes be denoted by d_A and d_B respectively. Figure 1(a)-1(b) plots the number of copies given to node B versus d_A for different values of l . For $l = 4$, the node which is closer to the destination gets most of the copies while for $l = 20$, most of the times, nearly half of the copies are given away to node B . This observation suggests that the optimal policy behaves differently for different values of l . (Note that node

¹The number of potential relays is equal to the number of nodes which do not have a copy of the packet. This number is upper bounded by the total number of nodes, M . Since the number of potential relays is unknown at a given time, we use the upper bound on this value.

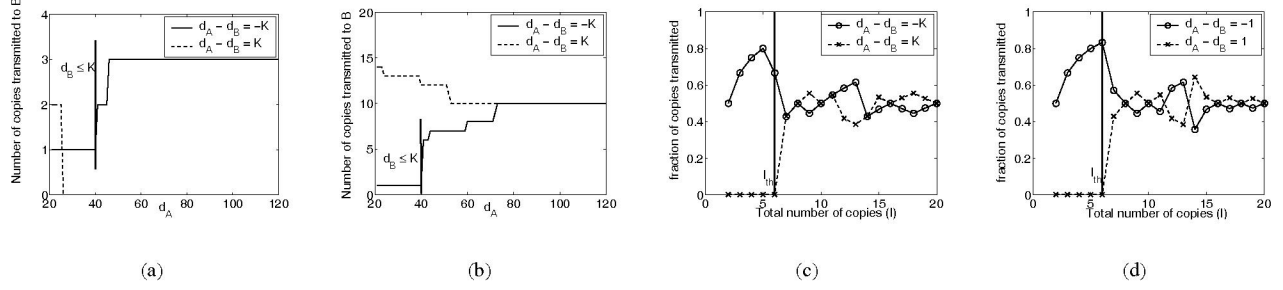


Fig. 1. Studying the optimal spraying policy for Spray and Wait. Network Parameters: $N = 150 \times 150$, $M = 40$, $K = 20$. (a) Number of copies given to node B as a function of d_A for $l = 4$. (b) Number of copies given to node B as a function of d_A for $l = 20$. (c) Proportion of copies given to node B as a function of l for $d_A = 75$. (d) Proportion of copies given to node B as a function of l for $d_A = 75$.




Fig. 2. Comparison of the expected end-to-end delay performance of binary spraying, the optimal policy and the proposed heuristic. Network parameters: $N = 150 \times 150$, $M = 40$, $K = 20$.




Fig. 3. Comparison of the expected end-to-end delay performance of binary spraying, the optimal policy and the proposed heuristic. Network parameters: $N = 150 \times 150$, $M = 40$, $L = 5$.

B gets only one copy when it is within the transmission range of the destination because the packet will be delivered at the next transmission opportunity.)

To study the behavior of the optimal policy as l changes, we plot the proportion of copies given to node B as a function of l for different values of $d_A - d_B$ in Figures 1(c)-1(d). In all the cases, there exists a threshold for l below which most of the copies are kept by the node closer to the destination and above which the copy splitting is more or less half and half. We label this threshold as l_{th} .

Based on the above observation, we propose a simple heuristic to distribute copies. (i) If l is less than l_{th} and node A is closer to the destination, then node B is not given any of the copies. (ii) If l is less than l_{th} and node B is closer to the destination, then node B is given $l-1$ copies. (iii) If l is greater than l_{th} , then node B is given half of the copies. Figures 2-3 compare the performance of the optimal policy, the proposed heuristic and binary spraying for different network parameters.

It is easy to see that the proposed heuristic performs very close to the optimal and has a better performance than binary spraying.

B. Spray and Focus

In Spray and Focus, each relay node performs utility based forwarding towards the destination. First, we derive the value of $E[D(d, 1)]$ to initialize the dynamic program which is used to solve Equation (5).

Lemma 4.1: $E[D(d, 1)]$ can be derived by solving the following set of non linear equations:

$$E[D(d, 1)] = \frac{1}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} + \frac{\frac{M}{E[M]}}{\frac{1}{E[M(d)]} + \frac{M}{E[M]}} \sum_{d_2} P(d_2 | d) E[D(\min(d, d_2), 1)]. \quad (6)$$

Proof: In the future time slots either of the following two events can happen first: (i) The node meets the destination and delivers the packet. This time duration is exponentially distributed with mean $E[M(d)]$. (ii) The node meets a potential relay node. This time duration is exponentially distributed with mean $\frac{E[M]}{M}$. Let the relay node be at a distance d_2 from the destination. Then if $d_2 < d$, then the relay node is closer to the destination and it will be given the copy of the packet. The additional time it will take to deliver the packet will be equal to $E[D(d_2, 1)]$. But if $d_2 \geq d$, the original node will retain the copy and the additional time it will take to deliver the packet is still equal to $E[D(d, 1)]$. The value of $P(d_2 | d)$ was derived at the end of Section III. \square

A particular value of $E[D(d, 1)]$ depends only on those values of $E[D(\hat{d}, 1)]$ for which $\hat{d} \leq d$. Hence, a dynamic program can be used to solve Equation (6).

Now, we study the optimal spraying policy obtained by solving Equation (5) after substituting the value of $E[D(d, 1)]$ derived in Lemma 4.1. Figure 4(a)-4(b) plots the number of copies given to node B versus d_A for different values of l . The curves show that most of the times, nearly half of the copies are handed over to node B irrespective of the value of l . To confirm this observation, we plot the proportion of copies given to node B as a function of l for different values of $d_A - d_B$ in Figures 4(c)-4(d). For all the cases, nearly half of

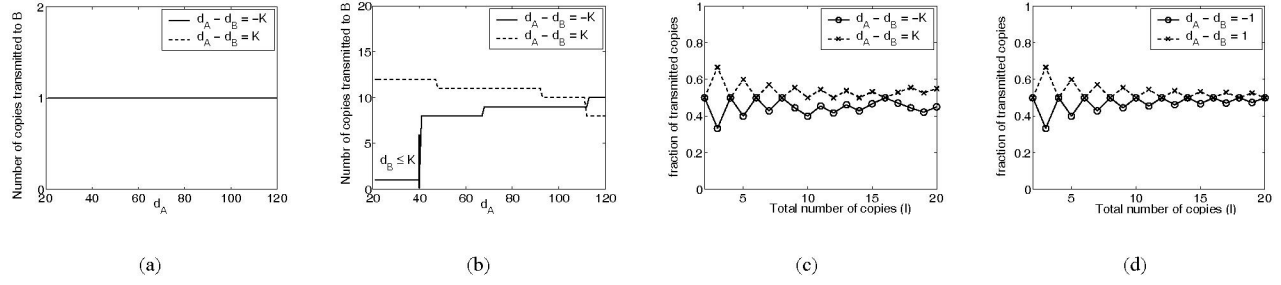


Fig. 4. Studying the optimal spraying policy for Spray and Focus. Network Parameters: $N = 150 \times 150$, $M = 40$, $K = 20$. (a) Number of copies given to node B as a function of d_A for $l = 2$. (b) Number of copies given to node B as a function of d_A for $l = 20$. (c) Proportion of copies given to node B as a function of l for $d_A = 75$. (d) Proportion of copies given to node B as a function of l for $d_A = 75$.




Fig. 5. Comparison of the expected end-to-end delay performance of binary spraying and the optimal policy. Network parameters: $N = 150 \times 150$, $M = 40$, $K = 20$.




Fig. 6. Comparison of the expected end-to-end delay performance of binary spraying and the optimal policy. Network parameters: $N = 150 \times 150$, $M = 40$, $L = 5$.

the copies are handed over to node B . This suggests that binary spraying should perform close to the optimal policy. Figures 5-6 compare the performance of binary spraying with the optimal policy for different network parameters. These figures show that binary spraying has near optimal performance for Spray and Focus. The near optimal performance of binary spraying is explained by the following two observations: (i) If a node distributes its copies to bad nodes (nodes which have a higher expected delivery delay), it still has its own copy which it can give to a good node whenever it meets one. (ii) Moreover, a bad node will have a chance to give up its copy to good nodes later in the future. Thus, spraying copies as fast as possible will achieve a good delay performance for Spray and Focus.

V. CONCLUSIONS AND FUTURE WORK

This paper presents a methodology to find the optimal way to distribute copies amongst the potential relays for spraying based schemes. As case studies, we find the optimal spraying

policies for Spray and Wait and Spray and Focus. We study the optimal policy to infer simple heuristics which achieve expected delays very close to the optimal. In ongoing work, we are trying to generalize the methodology so that it works for any mobility model as well as any other utility function. Additionally, we are trying to find an analytical plug and play expression for l_{th} in terms of the network parameters.

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 zebrant," in *Proceedings of ACM ASPLOS*, 2002.
- [2] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedevski, J. Pal, R. Patra, S. Surana, and K. Fall, "The case for technology in developing regions," 2005.
- [3] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Proceedings of ACM SIGCOMM*, Aug. 2004.
- [4] "Disruption tolerant networking," <http://www.darpa.mil/ato/solicit/DTN/>.
- [5] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, and K. Scott, "Delay-tolerant networking: an approach to interplanetary internet," *IEEE Communications Magazine*, vol. 41, 2003.
- [6] J. Ott and D. Kutscher, "A disconnection-tolerant transport for drive-thru internet environments," in *Proceedings of IEEE INFOCOM*, 2005.
- [7] C. Perkins, *Ad Hoc Networking*, 1st ed. Addison-Wesley, 2001.
- [8] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Duke University, Tech. Rep. CS-200006, Apr. 2000.
- [9] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic routing in intermittently connected networks," *SIGMOBILE Mobile Computing and Communication Review*, vol. 7, no. 3, 2003.
- [10] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," in *Proceedings of IFIP Networking*, 2006.
- [11] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and wait: Efficient routing in intermittently connected mobile networks," in *Proceedings of ACM SIGCOMM workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [12] T. Small and Z. Haas, "Resource and performance tradeoffs in delay-tolerant wireless networks," in *Proceedings of ACM SIGCOMM workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [13] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: The multi-copy case," *submitted to IEEE Transactions on Networking*.
- [14] Y. Wang, S. Jain, M. Martonosi, and K. Fall, "Erasure coding based routing for opportunistic networks," in *Proceedings of ACM SIGCOMM workshop on Delay Tolerant Networking (WDTN)*, 2005.
- [15] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications*. Birkhauser, Boston, 1998.
- [16] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Single-copy routing in intermittently connected mobile networks," in *Proceedings of IEEE SECON*, 2004.