

On the Benefits of Random Linear Coding for Unicast Applications in Disruption Tolerant Networks

Xiaolan Zhang, Giovanni Neglia[§], Jim Kurose, Don Towsley
University of Massachusetts
{ellenz,kurose,towsley@cs.umass.edu}

[§]Università degli Studi di Palermo
giovanni.neglia@tti.unipa.it

Abstract— In this paper, we investigate the benefits of using a form of network coding known as Random Linear Coding (RLC) for unicast communications in a mobile Disruption Tolerant Network (DTN) under epidemic routing. Under RLC, DTN nodes store and then forward random linear combinations of packets as they encounter other DTN nodes. We first consider RLC applied to a single block of K packets where (a) all K packets have the same source and destination, (b) the K packets have different sources but a common destination and (c) the K packets each have a different source/destination pair; we also consider the case where blocks of K packets arrive according to a Poisson bulk arrival process. Our performance metric of interest is the delay until the last packet in a block is delivered. We show that for the single block case, when bandwidth is constrained, applying RLC over packets destined to the same node achieves (with high probability) the minimum delay to deliver the block of data. We find through simulation that the benefit over non-network-coded packet forwarding increases further when buffer space within DTN nodes is limited. For the case of multiple blocks, our simulations show that RLC offers only slight improvement over the non-coded scenario when only bandwidth is constrained, but more significant benefits when both bandwidth and buffers are constrained. We remark that when the network is relatively loaded, RLC achieves improvements over non-coding scheme only if the spreading of the information is appropriately controlled.

I. INTRODUCTION

Epidemic routing ([13], [11], [12], [8], [16]) has been proposed for routing in mobile disruption tolerant networks (DTNs) in which there may not be a contemporaneous path from source to destination. Epidemic routing adopts a so-called “store-carry-forward” paradigm – a node receiving a packet buffers and carries that packet as it moves, passing the packet on to new nodes that it encounters. Analogous to the spread of infectious diseases, each time a packet-carrying node encounters a new node that does not have a copy of that packet, the carrier is said to *infect* this new node by passing on a packet copy; newly infected nodes, in turn, behave similarly. The destination receives the packet when it first meets an infected node.

Random Linear Coding (RLC) is a form of network coding [2] where each network node, rather than forwarding packets unchanged along the path from source-to-destination, can forward random linear combination of the data it has received.

In this paper, we investigate the use of RLC in epidemic routing for unicast applications in mobile DTNs through simulation. In this case there are different possible ways to

combine packets: each nodes can combine all the packets in its buffer, or only the packets destined to the same destination, or only the packets belonging to the same flow (i.e., same source-destination pair). We first consider these three possibilities in the simpler case of a single block of K packets. We then consider the case where blocks of K packets arrive according to a Poisson bulk arrival process. Our performance metric of interest is the delay until the last packet in a block is delivered. We show that for the single block case, when bandwidth is constrained, applying RLC over packets destined to the same node achieves (with high probability) the minimum delay to deliver a block of data. We find through simulation that this benefit increases further when buffer space within the DTN nodes is limited. For the case of multiple source/destination pairs, our simulations show that RLC offers only slight improvement over the non-coded scenario when only bandwidth is constrained, but more significant benefits when both bandwidth and buffers are constrained. We also demonstrate the “price” to be paid for the improved performance is generally a higher number of epidemically-spread copies of data in the DTN.

Several previous research efforts have applied *source-based* (i.e., non-network-coded) erasure codes to DTNs. [14] proposes erasure-coding-based routing for *opportunistic networks*, where DTN nodes operate without prior knowledge of node mobility patterns. For the case that a DTN has prior knowledge about paths and their loss behavior, [10] considers how to allocate the source-erasure-coded blocks to these paths. [6] proposes the usage of decentralized erasure codes to spread and store distributed data in a large scale sensor network. We note that this is similar to [14], in the sense that original packets traverse two hops to the final destination.

RLC has also been previously applied to networked scenarios including P2P content distribution ([7]), multicast application([3]), gossip protocol ([5], [4]) and distributed storage ([4], [1]). To our knowledge, the only work applying network coding in a DTN is [15], where the authors consider *broadcast* data delivery using RLC; our focus here is on using RLC for unicast delivery in a DTN.

The remainder of this paper is structured as follows. We introduce the network model and the forwarding schemes in Section II. The simulation setting is described in Section II. Section IV studies the benefit of RLC over non-coded scheme

for the scenario where there is a single generation of packets in the network. Section V extends the study to multiple generation case. Finally, Section VI summarizes the paper and discusses future work.

II. NETWORK MODEL AND FORWARDING SCHEMES

We consider unicast communications (i.e. each messages has a single node as destination) in a network consisting of N nodes moving according to a mobility model (discussed shortly) within a closed region. Each node has a fixed limited transmission range, such that the network is sparse and therefore disconnected. When two nodes come within transmission range of each other (i.e., they meet), they first figure out if the other has some useful information and, if any, they try to exchange it. We detail this process with reference to the two mechanisms we are going to compare: traditional non-coded packet-forwarding and RLC-forwarding.

Non-coded forwarding: When two nodes meet, each of them randomly selects one or more packets, depending on the bandwidth, among the packets that the other node does not have, and forwards them to the other node. We refer to this as the **random** selection scheme. We also consider a **RR_random** scheme in which the packet's source node chooses a packet to forward in round-robin manner, while intermediate nodes use random selection. Our intuition is that RR selection will help to speed up the propagation of initial copies of each packet.

Random Linear Coding based forwarding: RLC is applied to a finite set of K packets, called *generation*. Each packet is viewed as a d dimensional vector over a finite field, F_q of size q . We denote by $m_i \in F_q^d, i = 1, 2, \dots, K$ the K packets. A linear combination of the K packets is:

$$f_i = \sum_{i=1}^K \alpha_i m_i, \alpha_i \in F_q.$$

Addition and multiplication are over F_q . Initially, the source node(s) carries the original packets (a linear combination with special coefficients $\alpha_i = 1, \alpha_j = 0, j \neq i$). If a node carries r independent linear combinations, we say that the rank of the node is r . When a node with rank r (with r linear combinations f_1, \dots, f_r) meets another node, it generates a random linear combination of the current combinations, by selecting random coefficients $\beta_1, \dots, \beta_r \in F_q$, and generates: $f_{new} = \sum_{j=1}^r \beta_j f_j$. Obviously, f_{new} is a linear combination of the original K packets. This new combination, along with the coefficients in terms of the original packets, is forwarded to the other node. When a node (e.g., the destination) reaches rank K , it can decode the original K packets through matrix inversion.

Notice that RLC based scheme incurs storage overhead for storing coefficients for each combinations, and it also requires more computation to check if one node has useful information for the other and to decode the combinations.

Under epidemic routing, when a packet is delivered to the destination, a recovery scheme can be used to delete obsolete copies of the packet from the network [9]. We will use VACCINE recovery scheme throughout this paper. Under VACCINE, when a packet is first delivered, an *antipacket*

is generated and propagated through the network to delete buffered copies of this packet. To simplify analysis and simulation, we assume that the recovery scheme is not subject to bandwidth and buffer constraints. Under RLC scheme, when a generation is delivered to the destination, an antipacket for the generation is generated and propagated to delete buffered combinations of the generation.

We study the time to delivery a block of K packets when the packets are forwarded without any coding or when RLC is applied to the block of K packets. In particular, we define *block delivery delay* as the time from the arrival of the block in the network to the delivery of the whole block to the destination, denoted by D_{block} . Depending on the specific application, other metrics could be more meaningful, like the average time to deliver a packet of the block, or the average time to deliver a packet respecting the order. Note that D_{block} is the metric more favorable to RLC in the comparison. Another performance metric of interest is the average number of packet copies or combinations made within the network, as this is a measure of resources consumed (bandwidth, transmission power, buffering) within the DTN.

III. SIMULATION SETTING

We perform our simulation study using our own simulator. Rather than simulating a specific mobility model (e.g., the random waypoint or random direction mobility model), we directly simulate a pair-wise Poisson meeting process between two nodes. [8] has shown that under the random waypoint/direction models, the inter-meeting time between a pair of nodes follows a Poisson process when node velocity is relatively high compared to the region size, and the transmission range is relatively small. This simplification speeds up the simulation. We have also performed simulations using the actual mobility models and observe similar performance. Due to space constraints, these latter results are not presented here. For the results presented in this paper, we simulate a network of $N = 101$ nodes with a pair-wise meeting rate of $\beta = 0.0049$. We use a finite field of size $q = 701$.

IV. SINGLE GENERATION CASE

In this section, we focus on the simple setting in which there is a single generation of packets in the network. In particular we assume that K packets arrive at the same time in the network. We examine the following three scenarios:

- **SS_SD (Single Source/Single Destination):** in which data in K packets from a source are to be delivered to a single destination;
- **MS_SD (Multiple Source/Single Destination):** in which data in K packets from different sources are to be delivered to the same destination;
- **MS_MD (Multiple Source/Multiple Destination):** in which data in each of K packets (each from a different source) are to be delivered to a different destination.

A. Benefit of coding under bandwidth constraints

We first consider the case when bandwidth is constrained, i.e., when two nodes meet, they can send a maximum of b packets in each direction. We assume for now that mobile nodes have sufficient buffer space to store all packets.

Claim 1: If there is a single block of packets in the network, for the SS_SD and MS_SD case, RLC achieves the minimum D_{max} with high probability.




Fig. 1. Random graph representing the contacts between nodes

Consider a random multigraph constructed as follows (Fig.1): there are N vertices, each corresponding to one mobile node, and for each contact between a pair of nodes that can exchange b packets in each direction, add b directed edges in each direction between the corresponding vertices. Edges are labeled with the time that the contact occurs. A *path* in such a network is a path in the graph where the successive edges have increasing timestamps. A set of paths are *independent* if they do not share edges.

For the SS_SD case, where the source node initially has K packets to send at $t = 0$, the time to deliver these K packets cannot be smaller than the time when there are K independent paths from the source to the destination. Similarly, for the MS_SD case, the delivery time cannot be smaller than the time to have K independent paths from the K source nodes to the destination.

For non-coding scheme, this minimum delay is hard to achieve, since a node has only local information (i.e., is not aware of packets transfers happening between other nodes), it is likely that the nodes along some of these K paths forward packets that other paths are propagating. Under RLC, due to the randomization, with high probability, the first K combinations delivered to the destination by these K paths are independent (therefore the original K packets can be decoded achieving the minimum delay), as illustrated below.

Fig.1 illustrates this idea using an example of a 4-node network. Assume that node 1 generates two packets m_1, m_2 destined to node 4 at $t = 0$. Without applying network coding, node 1 forwards m_1, m_2 to node 2, and one of the packets (say m_1) to node 3. When nodes 2 and 4 meet at $t = 10.2$, node 2 randomly selects a packet and delivers to node 4 (given that the node has no global knowledge of past and future contacts for other nodes). With probability 0.5, packet m_2 is forwarded to node 4. As a result, when node 3 meets node 4 at $t = 23$, it has no useful information for node 4. On the other hand, if

RLC is used, suppose source node 1 forwards random linear combination c_1, c_2 to node 2, and c_3 to node 3. On meeting node 4, node 2 generates a random linear combination c_{12} of c_1, c_2 and forwards it. As long as c_{12} and c_3 are independent (with probability $1 - 1/q$, if the coding is over finite field F_q), node 4 can decode the two original packets after node 3 delivers c_3 at time $t = 23$.

As a quantitative analysis of delivery delay is difficult due to the random nature of the contacts, and the large size of the networks in which we are interested, we use simulation to quantify the performance gain of RLC scheme.

We first highlight several characteristics of RLC compared to a non-coded approach. Fig.3 depicts the total number of packet copies (for the non-coding scheme) or combinations (for RLC scheme) in the entire network as a function of time for SS_SD under $N = 100, K = 10$ in a particular run. The graph shows that RLC allows faster propagation of information. On the other hand, it also shows that RLC incurs more copies being made in the network. There are two factors causing more transmissions made under RLC: first, RLC allows faster propagation of information, secondly, under RLC scheme the recovery process starts only when the whole generation is delivered (much slower than under non-coding approach, where recovery process for individual packet starts immediately when it is delivered).







Fig. 3. RLC scheme achieves faster propagation

We next explore the relative benefit of RLC with respect to the non-coded case under varying bandwidth constraints. Fig.2(a) plots the average D_{block} for SS_SD with $K = 10$ under varying bandwidth constraints. (We note that the average D_{block} reported throughout Section IV are the average value from 50 different simulation runs). The figure shows that RLC achieves lower D_{block} than both random and RR_random schemes. Furthermore, the relative benefit of RLC increases as bandwidth decreases.

Fig.2(b) shows the sensitivity of performance to block size, plotting the average D_{max} for the SS_SD case with varying block size and a bandwidth constraint of $b = 1$ (i.e., on every contact, only one packet can be sent in each direction. For the remainder of this paper, this is the default bandwidth constraint used in our simulation results). From Fig.2(c), we observe that as the block size increases, the relative benefit of RLC over non-coding scheme decreases.

(a) D_{max} under varying bandwidth(b) D_{max} under different block size

(c) Benefit of RLC under different block size

Fig. 2. RLC benefit under SS_SD

Due to space constraint, our results for MS_SD and MS_MD case are not shown here. We note that the benefit achieved by RLC for the MS_SD case is smaller than for the SS_SD case. This is because here the K packets start to propagate from K different nodes, the effect of relay nodes choosing the wrong packets to forward becomes less significant. For MS_MD case, we find that with only bandwidth constraints, RLC performs worse than the non-coding scheme since RLC forces every destination node to receive K independent combinations to decode the one single packet destined to it.

B. Benefit of Coding under Bandwidth and Buffer Constraints

Thus far, we have assumed that nodes have unlimited buffer capacity. In this section, we assume that the relay nodes can store at most B ($B < K$) packets or combinations; source and destination nodes are not subject to this constraint. For RLC, when a node receives a combination and its buffer is full, it randomly combines the new combination with an existing combination in the buffer and stores the result. For the non-coding case, a drophead scheme ([16]) is used which drops the packet that has resided in the buffer the longest when a new packet arrives and the buffer is full.

Fig.4(a) shows that, for the SS_SD case (with $K = 10$), as nodal buffer sizes decrease, the performance of RLC degenerates only slightly; while the performance of the non-coding schemes degrade quickly. However, this is achieved at the cost of more transmission made as shown in Fig.4(b). Notice that although under unconstrained buffer case, at most K linear combinations of a generation (of size K) are sent to each node, this is not the case under buffer constraint where a node can be repeatedly sent different combinations of a generation without increasing rank. For MS_SD case, we observe similar performance gains of RLC (not shown here). Under MS_MD ($K = 10$), where coding is applied to packets sent by different sources to different destinations, when the buffer is very constrained ($b = 1$ for this setting), RLC out-performs the non-coding random scheme (Fig.4(c)). The benefit of RLC under bandwidth constraints comes from the fact that random linear combinations of packets increases the tolerance to losses due to buffer overflow.

C. Controlling Transmission Power Consumption of RLC

We have seen that RLC is able to deliver a block of data, or collect multiple packets from different sources in the minimum


amount of time, at the “cost” of having more copies of packets present in the network, consuming more buffer space and more transmission power (to send these copies), as shown in Fig.3. In this section, we study the question of whether RLC can achieve a performance gain under the same transmission power consumption as a non-coded scheme.

To limit the number of copies made for a packet, we use a token-based scheme, extending the *spray and wait* scheme proposed in [12], [11]. Under this scheme, every new packet generated at the source is assigned a certain number of tokens (which we refer to as per-packet token number). When the packet is copied to another node, the token number is decreased by one, and half of remaining tokens are assigned to the new copy. When a packet copy has only a single token remaining, it can only be forwarded to the destination. The total number of copies of a packet is thus bounded by the initial token number. We note that this scheme can be improved by allowing two nodes carrying copies of the same packet to average their token numbers when they meet. We extend the notion of tokens to the RLC scheme by associating a token number with the generation, which equals to the product of the number of packets in the generation and the per-packet token number. Instead of splitting tokens in equal halves when making a new copy and when two nodes meet, the token number for a generation are allocated to two nodes in proportion to their ranks (i.e., the *amount* of information the nodes carry about the generation).


We run simulations for SS_SD case with $K = 10$ with different per-packet token numbers. We notice that under the same per-packet token number, RLC makes more copies because the random mixing enables nodes to have information to exchange more often. Nevertheless, RLC scheme achieves better transmission/delay tradeoff than non coding scheme, as shown in Fig.5. The figure plots for SS_SD case with $K = 10$, the number of transmissions versus delay tradeoff achieved when the per-packet token limits are varied between 4 and 90, and unlimited. It shows that under similar transmission consumption, RLC achieves smaller delivery delay.

V. MULTIPLE GENERATION CASE


In this section, we investigate whether RLC is beneficial in a more complex scenario where many generations are present at the same time in the network. In particular we consider multiple asynchronous continuous unicast flows in the network



(a) D_{block} for SS_SD under varying buffer



(b) Number of transmissions made for SS_SD



(c) D_{block} for MS_MD under varying buffer size

Fig. 4. Bandwidth and buffer constraint

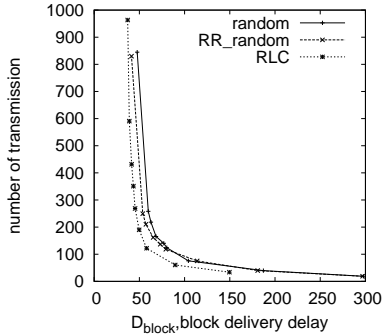


Fig. 5. Transmission power vs delay trade-off achieved with different token number

and coding within the same flow. In what follows, we first introduce the traffic process and scheduling schemes, and then present the results for the two scenarios previously considered: when only bandwidth is constrained, and when both bandwidth and buffer are constrained.

A. Settings: Traffic Process and Scheduling Schemes

We assume there are N flows in the network, with each node being the source of one flow and the destination of one other flow. Each source node generates independently a block of $K = 10$ packets according to Poisson process with rate λ . Therefore the total packet arrival rate to the network is given by $NK\lambda$. We only consider applying RLC to packets belonging to the same block, i.e. each block forms a generation.



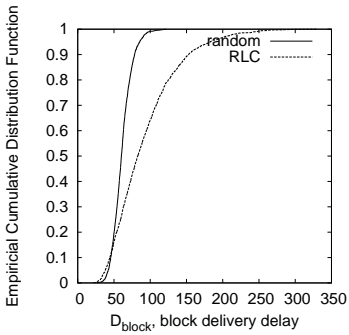
As our focus is not scheme design, but rather the understanding of the benefit of RLC, we adopt simple randomized scheduling schemes. For non-coding scheme, when a node meets another node, it randomly selects a packet from the set of packets that it carries while the other node does not have, and forwards it. For RLC scheme, the node first randomly chooses a generation from the set of generations that it carries which have some useful information for the other node, and then generates a random linear combination for this generation to forward. For both cases, priorities are given to the packets/generations destined to the other node; furthermore, among such packets/generations, those originated from the node are served first.

B. Benefit of Coding under Bandwidth Constraint

We have seen that under bandwidth constraint, for one single generation, RLC achieves lower delay than non-coding scheme, because RLC is able to take advantage of a larger number of contacts. We perform simulation studies under varying block arrival rate with bandwidth constraint $b = 1$. We observe that RLC only has benefit when the traffic rate is low; and performs worse than non-coding scheme when the traffic rate is high, as shown in Fig.6(a) which plots the empirical cumulative distribution function (CDF) of D_{block} under $\lambda = 0.00045$.

The reason is two-fold. First, for non-coding scheme, when the arrival rate λ increases, the number of different packets in the network increases and it is more likely that two nodes have some useful information to exchange when they meet, therefore the gain of RLC is smaller. Secondly, as we have shown in Fig.3, RLC generates more transmissions for each generation. When the block arrival rate is high and there are many generations in the network at the same time, these different generations start competing for the bandwidth. In fact an optimal scheduling should favor a new generation over an old generation that has a larger number of combinations spreaded in the network (and with high probability of being already delivered. But currently implemented random scheduling scheme does not consider this optimization.

The tradeoff between number of transmissions and average delay shown in Fig.5 suggests a way to deal with this resource contention problem. The figure shows that RLC can achieve the same delay as non-coding with a significantly lower number of transmissions (left part of the curve), so we expect significant benefit by appropriately limiting copies made for a generation. Fig.6(b) shows that this is the case. Fig.6(b) plots the average D_{block} achieved for RLC and random schemes under block arrival rate of $\lambda = 0.00045$, when the per-packet token limit is varied between 20 and 100. In particular there is an optimal token limit value for RLC scheme, between 40 and 50 token. For higher values, the contention degrades the performance, while for lower values some useful meeting cannot be exploited because all the tokens have been consumed. For non-coding scheme under this arrival rate, the contention is not significant and the reduction of the number of tokens incurs a larger delay (note that we observe that under a higher block arrival rate, non-coding scheme also benefits from limiting the



(a) Empirical CDF of D_{block} under $\lambda = 0.00045$

(b) D_{block} under different token limit

(c) D_{block} under different token limit, $B = 5$

Fig. 6. Block delivery delay under multiple generation case

number of copies). We can estimate the maximum number of transmissions that can be made for each packet as the ratio between the total bandwidth available in the networks, $\frac{1}{2}N(N-1)\beta$, and the total arrival rate, $NK\lambda$. For the specific setting considered here, this value is equal to 50, which is close to the optimal token value found by simulation.

C. Bandwidth and Buffer Constrained Case

Section IV-B has shown that for a single generation case, RLC is especially useful when buffer is constrained. We consider the scenario when there are multiple generations in the network in this section. As usual, we assume that each node has limited buffer for storing relay packets, but unlimited buffer for storing its own source packets. Since the source node always stores a packet until it is known to be delivered, there is no packet loss. When a node receives a combination and its buffer is full, it first selects randomly one generation from the generations in its buffer that have the highest rank. If the new combination is for the chosen generation, the new combination is combined with an existing combination within that generation. Otherwise, the node compresses the matrix of the selected generation by one, and the new combination is inserted to its generation's matrix.

When both bandwidth and buffer are constrained, limiting number of transmissions made for a generation becomes even more important for RLC scheme. As Fig.4(b) in Section IV-B shows, under a single generation case, RLC scheme makes much more transmissions than non-coding scheme. Therefore, when there are multiple generations in the network, resource contention is even greater than when buffer is not constrained. We expect that using token scheme allows to allocate bandwidth and buffer space more evenly among different generations. We simulate the case of block arrival rate of $\lambda = 0.00045$, and every node only store $B = 5$ relay packets (combinations) under various token limits. As Fig.6(c) shows, the RLC achieves lower block delivery delay than non coding scheme, reducing the delay by about 22.5%.

VI. SUMMARY

We have studied the benefits of applying RLC to unicast application in mobile DTN in this paper. For the case where there is a single generation in the network, we found that RLC applied to a block of data destined to the same destination achieves minimum block delay with high probability. Larger

gain is achieved by RLC scheme when furthermore buffer space is constrained. Although RLC scheme makes more transmissions, by using token limit scheme, RLC scheme can achieve better transmission power/delay tradeoff than non-coding approach. When there are multiple generations in the network, under appropriately chosen token limit, RLC scheme achieves slight gain over non-coding scheme under only bandwidth constraint, and a significant gain when nodal buffer is also constrained.

REFERENCES

- [1] S. Acedanski, S. Deb, M. Medard, and R. Koetter. How good is random linear coding based distributed networked storage. In *Netcod*, 2005.
- [2] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Trans. on Information Theory*, 46:1204–1216.
- [3] P. A. Chou, Y. Wu, and K. Jain. Practical network coding. *Allerton Conference on Communication, Control, and Computing*, October.
- [4] S. Deb, C. Choute, M. Medard, and R. Koetter. Data harvesting: A random coding approach to rapid dissemination and efficient storage of data. preprint.
- [5] S. Deb and M. Medard. Algebraic gossip: A network coding approach to optimal multiple rumor mongering. 2004. Proc. Allerton.
- [6] A. G. Dimakis, V. Prabhakaran, and K. Ramchandran. Ubiquitous access to distributed data in large-scale sensor networks through decentralized erasure codes. In *Symposium on Information Processing in Sensor Networks (IPSN '05)*.
- [7] C. Gkantsidis and P. Rodriguez. Network coding for large scale content distribution. *Infocom*, 2005.
- [8] R. Groenevelt, P. Nain, and G. Koole. The message delay in mobile ad hoc networks. In *Performance*, October 2005.
- [9] Z. J. Haas and T. Small. A new networking model for biological applications of ad hoc sensor networks. to appear in *IEEE/ACM Transactions on Networking*.
- [10] Sushant Jain, Michael Demmer, Rabin Patra, and Kevin Fall. Using redundancy to cope with failures in a delay tolerant network. In *SIGCOMM*, 2005.
- [11] T. Small and Z. J. Haas. Resource and performance tradeoffs in delay-tolerant wireless networks. In *ACM workshop on Delay Tolerant Networking*, 2005.
- [12] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In *ACM workshop on Delay-tolerant networking*, 2005.
- [13] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-200006, Duke University, April 2000.
- [14] Y. Wang, S. Jain, M. Martonosi, and K. Fall. Erasure-coding based routing for opportunistic networks. *ACM workshop on Delay Tolerant Networking*, 2005.
- [15] J. Widmer and J.-Y. Le Boudec. Network coding for efficient communication in extreme networks. *ACM workshop on Delay Tolerant Networking*, 2005.
- [16] X. Zhang, G. Neglia, J. Kurose, and D. Towsley. Performance modeling of epidemic routing. Technical Report 2005-44, UMASS Computer Science. http://gaia.cs.umass.edu/pub/Zhang05_epidemic_TR.pdf.