Time Dependent Message Spraying for Routing in Intermittently Connected Networks

Eyuphan Bulut, Zijian Wang and Boleslaw K. Szymanski
Department of Computer Science and Center for Pervasive Computing and Networking
Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, USA
{bulute, wangz, szymansk}@cs.rpi.edu

Abstract—**Intermittently connected mobile networks, also called Delay Tolerant Networks (DTN) are wireless networks in which at any given time instance, the probability of having a complete path from a source to destination is low. Several routing algorithms are proposed for such networks based on flooding and erasure coding techniques. Since flooding based schemes suffer from huge overhead of bandwidth and energy consumption due to redundant transmissions, controlled flooding algorithms which use fixed number of copies for each message have been developed.**

Although a DTN is delay tolerant by definition, there still may be a required upper bound on the delay for the delivery of messages. In this paper, we propose a novel spraying algorithm in which the number of message copies in the network depends on the urgency of meeting the expected delivery delay for that message. The main objective of this algorithm is to give a chance to early delivery with small number of copies in existence, consequently decreasing the average number of copies sprayed in the network. We derive the formula for the optimum borders of periods for spraying for 2-period and 3-period variants of our algorithm. We also present the simulations of the method and compare their results with the analytical ones and observe the good match between them. Furthermore, we demonstrate that time dependent spraying algorithm provides a significant decrease in average copy count per message while preserving the percentage of the messages delivered before the upper bound of the acceptable delay expires.

I. Introduction

Intermittently connected mobile networks also referred to as Delay Tolerant Networks (DTNs) are wireless networks in which at any given time instance, the probability that there is an end-to-end path from a source to destination is low. There are many examples of such networks in real life including wildlife tracking sensor networks [1], military networks [2] and vehicular ad hoc networks [4]. Since the standard routing algorithms assume that the network is connected most of the time, they fail in routing of packets in DTNs.

Routing algorithms for DTNs need to carefully consider the disconnectivity of the network. Hence, in recent years, new algorithms using buffering and contact time schedules have been proposed. Since most of the nodes in a DTN are mobile, the connectivity of the network is maintained only when nodes come into the transmission ranges of each other. If a node has a message copy but it is not connected to another node, it stores the message until an appropriate communication opportunity arises. The important considerations in such a design are (i) the number of copies that are distributed to the network for each message, and (ii) the selection of nodes to which the message is replicated.

In this paper, we study how to distribute the copies of a message among the potential relay nodes in such a way that the predefined percentage of all messages meets the given limit for delay of delivery with the minimum number of copies used. Unlike than the previous algorithms, we propose a time dependent copying scheme which basically considers the time remaining to the given limit for the delay of delivery.

The idea of our scheme is as follows. We first spray some number of copies smaller than the necessary to guarantee that the predefined percentage of all messages meets the given limit for delay of delivery of the message to the destination. If the delivery does not happen for some time, then we spray some additional copies of the message to increase the probability of its delivery. Consequently, we will benefit from early delivery with less number of copies, if it happens, so the average number of copies used by each message is minimized.

The remaining of the paper is organized as follows. In Section II we present the previous work done on this topic and discuss some basic mobility assisted routing concepts. We also differentiate our algorithm from others. In Section III we describe our algorithm in detail and provide its analysis for its different variants. In Section IV, we present evaluation of the performance of the proposed scheme using simulations and demonstrate the achieved improvements. We also compare the results of our analysis with the simulation results. Finally, we offer conclusion and outline the future work in Section V.

II. Related Work

Routing algorithms for delay tolerant networks are generally classified as either replication based or coding based [13]. In replication based algorithms, multiple or a single copy of the message is generated and distributed to other nodes (often referred to as relays) in the network. Then, any of these nodes, independently of others, try to deliver the message copy to the destination. In coding based algorithms, a message is converted into a large set of code blocks such that any sufficiently large subset of these blocks can be used to reconstruct the original message. As a result, a constant overhead is maintained and the network is made more robust against the packet drops when the congestion arises. However, these algorithms introduce an overhead of an extra work needed for coding, forwarding and reconstructing code blocks.

Epidemic Routing [3] is an approach used by the replication based routing algorithms. Basically, in each contact between any two nodes, the nodes exchange their data so that they both have the same copies. As a result, the fastest spread of copies is achieved yielding the optimum delivery time. However, the main problem of this approach is the overhead incurred in bandwidth usage, buffer space required and energy consumed by the greedy copying and storing of messages. Hence, this approach is inappropriate for resource constrained networks. To address this weakness of epidemic routing, the algorithms with controlled replication or spraying have been proposed [5], [6], [7], [14]. In these algorithms, only a small number of copies are distributed to other nodes and each copy is delivered to the destination independently of others. Of course, such approach limits the aforementioned overhead and resources are efficiently used.

The replication based schemes with controlled replication differ from each other in terms of assumptions about the network. Some of them assume that the trajectories of the mobile devices are known while some others assume that the contact times and durations of nodes are known. There are also some algorithms which assume zero knowledge about the network. The algorithms which fall in this last category seem to be the most relevant to applications because in most of the examples of delay tolerant networks from real life, neither the contact times nor the trajectories are known for certain. Consider the difficulty of acquiring such information in a wild life tracking application where the nodes are attached to animals that move unpredictably.

The algorithms which assume zero knowledge about the network include [9], MaxProb [12], SCAR [11] and Spray and Wait [8]. In each of these algorithms limited number of copies are used to deliver a message. Yet, the process of choosing the nodes for placing new replications is different in each of them. In [9] and MaxProb each node carries its delivery probability which is updated in each contact with other nodes. If a node with a message copy meets another node that does not have the copy, it replicates the message to the contact node only if that node's delivery probability is higher than its own. A similar idea is used in SCAR. Each node maintains a utility function which defines the carrier quality in terms of reaching the destination. Then, each node tries to deliver its data in bundles to a number of neighboring nodes which have the highest carrier quality.

In [8] Spyropoulos et al. propose two different algorithms called Source Spray and Wait and Binary Spray and Wait, respectively. While in the former, only the source is capable of spraying copies to other nodes, in the latter all nodes having the copy of the message are also allowed to do so. In Binary Spray and Wait, when a node copies a message to another node, it also gives the right of copying the half of its remaining copy count to that node. This results in distributed and faster spraying compared to the source spraying, but once the spraying is done, the expected delivery delay is the same. The authors provide the expected delay of message delivery in these two algorithms in [14].

Although there are many algorithms utilizing the controlled flooding approach, the idea of copying dependent on the urgency of meeting the limit on delay for delivery of a message has not been used by any of them. To the best of our knowledge this idea is new and it helps to decrease the average number of copies generated in the network. We will describe the details of this idea in the next section.

While designing a routing algorithm for mobile network, an important issue that must be considered is the model of mobility of nodes in the network. Random direction, random walk and random waypoint mobility models are the most popular ones among those used by the previous routing algorithms in this field. Among these models, random direction model is considered more realistic than the others.

In a network with mobile nodes moving according to a mobility model, two concepts are introduced; expected hitting time (ET) and expected meeting time (EM). While ET of a node is defined as the expected time interval of being in contact (in the transmission range) with a stable node (most often the sink), EM of a node is defined as the expected time interval of being in contact with a nearby mobile node. These two parameters are specific to each mobility model and can be derived when the network parameters are known [10]. Then, the hitting time and the meeting time of a sample node in such a network are assumed to be exponentially distributed with mean ET and EM, respectively.

III. Time Dependent Spraying

In this section, we start with listing the assumptions of our model and then we provide the details of our routing scheme and its analysis.

We assume that there are M nodes randomly walking on a $\sqrt{N}$ x $\sqrt{N}$ 2D torus according to the random direction mobility model. Each node has a transmission range R and all nodes are identical. The buffer space in a node is assumed to be infinite (not crucial since we use less copies), and the communication between nodes is assumed to be perfectly separable, that is, any communicating pair of nodes do not interfere with any other simultaneous communication. To be consistent with previous research, by L we denote the number of copies distributed to the network.

In some studies, authors find out the minimum number of copies needed to achieve a given delay with predefined probability. As discussed previously, the optimal delay in a mobile network is obtained by epidemic routing in which there is a complete message exchange in every contact of any two nodes. Figure 1 shows the minimum number of copies (L_{min}) needed to achieve the expected delay which is $'a'$ times the optimal delay [14].

Given the mobility model, the expected time delivery by Spray and Wait algorithm is equal to [14]:

$$\sum_{i=1}^{L-1} \frac{EM}{M-i} + \frac{M-L}{M-1} EW$$

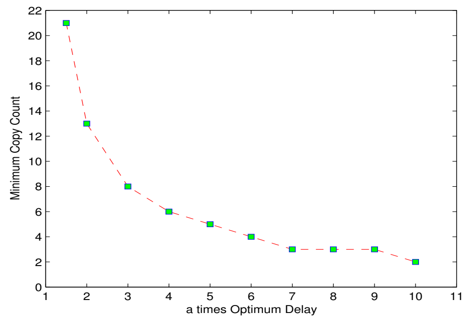

Fig. 1. Minimum L needed to meet the delay equal to $'a'$ times the optimum delay.

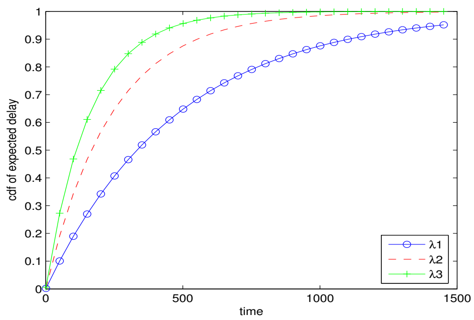

Fig. 2. The cumulative distribution function of probability of meeting the expected delay in Spray and Wait algorithm for different λ values, where $\lambda_1 > \lambda_2 > \lambda_3$

This formula assumes that in the first $L - 1$ contacts, the source node does not meet with the sink node and thus a wait phase is needed (probability of this happening is $\frac{M-L}{M-1}$). Here, EW is the expected duration of wait phase which is actually exponentially distributed with mean $\frac{EM}{L}$. Note that, when $M >> L$ (which we enforce to be satisfied by limiting permissible values of L), duration of spraying phase is much shorter than the duration of waiting phase, so that we can assume that the expected time of delivery in Spray and Wait algorithm is exponentially distributed with mean $\frac{EM}{L}$.

Figure 2 shows the cumulative distribution function of the expected delay of Spray and Wait algorithm for different L values. Clearly, when L increases, mean value $(1/\lambda)$ decreases and the expected delay shrinks.

Our contribution to the spray and wait idea is to control spray of packets to other nodes by the urgency of meeting the predefined delivery delay. More precisely, the algorithm starts with spraying the message copies to fewer nodes than the minimum L needed and then waits to see if the message is delivered for a certain period of time. When delivery does not happen, the algorithm increases the number of copies sprayed and again waits for delivery. This process repeats until the message is delivered or the time limit for delivery is reached. Hence, as the time remaining to the limit for delivery time

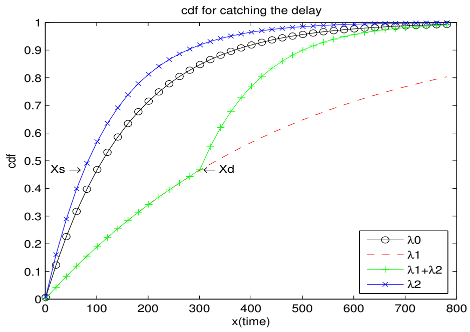

Fig. 3. The cumulative distribution function of delivery time of a message with spraying different number of copies in two different periods.

decreases and delivery has not happened[1], the number of nodes carrying the message copy increases. To the best of our knowledge, this idea has not been used by any of the previously published algorithms for DTN routing.

Consider the Figure 3. It summarizes what our algorithm wants to achieve. In this specific version of the algorithm, we allow two different spraying phases. The first one is done at the beginning and the second one is done at time x_d. The main objective of the algorithm is to attempt delivery with small number of copies and use the large number of copies only when this attempt is unsuccessful. With proper setting, the average number of copies sprayed in the network till delivery will be lower than in case of spraying all messages at the beginning.

To analyze the performance of our algorithm analytically, we need to derive two formulas, one for the average copy count used by the algorithm, and the second one for the cumulative distribution of the probability of meeting the limit for the deliver delay with mixed number of copies (and therefore mixed λ values). The goal is to achieve the same percentage of the messages delivered in the given limit on the time of delivery using fewer copies on average than the standard Spray and Wait algorithm uses.

In our scheme, term *period* refers to the time duration from the beginning of one spraying phase to the beginning of the next spraying phase. There may be multiple spray phases and the corresponding periods between them, each of different length. We start with the analysis of the two period case to find out the optimal period length and the corresponding copy counts of each.

1) Two Period Case: Since there are two periods until the limit for the time of delivery of a message to the destination is reached, the arising questions are how the time should be divided into the two periods and how many copies should be allowed in each. In other words, what should be the value of x_d in Figure 3 to minimize the average copy count of the algorithm execution?

[1]We assume that the destination acknowledges received messages using a broadcast to all nodes, thereby suppressing any spraying after the message delivery. Such acknowledgments are short and can be broadcast using more powerful radio that often is present at the destination node.

Let's assume that the standard Spray and Wait algorithm uses L copies (including the copy in the source node) of a message to achieve the probability $p_d \approx 1$ of delivery of the message by the deadline t_d. Let's further assume that the Two Period Delayed Spraying algorithm sprays L_1 copies to the network at the beginning of execution and additional $L_2 - L_1$ copies at time x_d, the beginning of the second period. Then, the cumulative distribution function of the probability of delivering the message at or below time x is:

$$cdf(x) = \begin{cases} 1 - e^{-\alpha L_1 x} & \text{if } x \leq x_d \\ 1 - e^{-\alpha L_2 (x - x_s)} & \text{if } x > x_d \end{cases}$$

where, $\alpha = 1/EM$ is an inverse of the expected meeting time of the nodes and x_s is the time interval by which the second exponential function of the above formula is delayed compared to the first one and define by equality of the both functions at point x_d, hence:

$$\begin{aligned} 1 - e^{-\alpha L_1 x_d} &= 1 - e^{-\alpha L_2 (x_d - x_s)} \\ x_s &= x_d \frac{L_2 - L_1}{L_2} \end{aligned}$$

The expected delivery ratio when L copies are used in the standard Spray and Wait algorithm are by definition $p_d = 1 - e^{-\alpha L t_d} \approx 1$. We have tested the success rates of meeting the deadline with the different number of copies, where the delivery rate and L is chosen from the values in Figure 1. We want to match these delivery rates by decreasing the average number of copies below L, the number of copies used in the Spray and Wait algorithm. Hence, the following inequality must be satisfied:

$$\begin{aligned} 1 - e^{-\alpha L_2 (t_d - x_s)} &\geq 1 - e^{-\alpha L t_d} \\ L_2 (t_d - x_d + x_d L_1 / L_2) &\geq L t_d \end{aligned}$$

We can use this inequality to bound x_d as $x_d \leq t_d \frac{L_2 - L}{L_2 - L_1}$. It is clear that to minimize the average copy count in the two period case with the given L_1 and L_2 values, x_d should be as large as possible, hence

$$x_d = t_d \frac{L_2 - L}{L_2 - L_1}$$

We want to minimize the average number of packets, $c_2(L_1, L_2)$ defined as:

$$\begin{aligned} c_2(L_1, L_2) &= L_1 (1 - e^{-\alpha L_1 x_d}) \\ &\quad + L_2 [e^{-\alpha L_1 x_d} - e^{-\alpha L t_d}] \\ &= L_1 + (L_2 - L_1) e^{-\alpha L_1 x_d} - L_2 e^{-\alpha L t_d} \\ &\approx L_1 + (L_2 - L_1) e^{-\alpha L_1 x_d}. \end{aligned}$$

Substituting x_d in the above, we get:

$$c_2(L_1, L_2) = L_1 + (L_2 - L_1) e^{-\alpha L_1 t_d \frac{L_2 - L}{L_2 - L_1}}$$

Taking derivative of c_2 in regard of L_2, and comparing it to zero, we obtain:

$$L_2 = L1 + \alpha L_1 t_d (L - L_1)$$

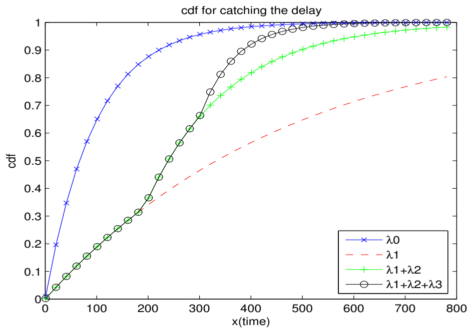

Fig. 4. The cumulative distribution function with spraying different number of copies in three different periods.

so $L_2 - L_1 = \alpha L_1 t_d (L - L_1)$ and therefore

$$c_2^*(L_1) = L_1 [1 + \alpha t_d (L - L_1) e^{-\alpha L_1 t_d + 1}].$$

Taking the derivative of the above function we can obtain the complicated formula on the optimal value of L_1^* as a function of L and t_d, and then taking the floor and ceiling compute the corresponding optimal values of L_2^* and again their floors and ceilings can be used to arrive at the result. A simpler method is to enumerate all integer values for L_1 from 1 to $L - 1$ and compute floors and ceilings of each corresponding optimal L_2 to compute the result.

2) Three Period Case: If there are three spray and wait periods, we need to find two different boundary points which separate these periods. Let x_{d1} and x_{d2} denote these boundary points, respectively. While the former stands at the boundary between the first and the second periods, the latter marks the boundary between the second and the third periods. The cumulative distribution function of the probability of delivering the message by the time x becomes:

$$cdf(x) = \begin{cases} 1 - e^{-\alpha L_1 x} & [0, x_{d1}] \\ 1 - e^{-\alpha L_2 (x - x_{s1})} & (x_{d1}, x_{d2}] \\ 1 - e^{-\alpha L_3 (x - x_{s2})} & (x_{d2}, x] \end{cases}$$

where x_{s1} and x_{s2} are the delays of the second and the third exponential functions compared to the first and can be easily computed as:

$$\begin{aligned} 1 - e^{-\alpha L_1 x_{d1}} &= 1 - e^{-\alpha L_2 (x_{d1} - x_{s1})} \\ x_{s1} &= x_{d1} \frac{L_2 - L_1}{L_2} \end{aligned}$$

and analogously

$$\begin{aligned} 1 - e^{-\alpha L_2 (x_{d2} - x_{s1})} &= 1 - e^{-\alpha L_3 (x_{d2} - x_{s2})} \\ x_{s2} &= x_{d2} \frac{L_3 - L_2}{L_3} + x_{d1} \frac{L_2 - L_1}{L_3}. \end{aligned}$$

Consider Figure 4 that illustrates our approach with three periods. Similar to the two period case, we want to achieve the same or higher delivery rate p_d within the given deadline t_d while minimizing the average number of copies used. That is, we need to satisfy the following inequality:

$$\begin{aligned} 1-e^{-\alpha L t_d} &\le 1-e^{-\alpha L_3(t_d-x_{s2})} \\ L_3(t_d-x_{s2}) &\ge L t_d \\ x_{d2}(L_3-L_2)+x_{d1}(L_2-L_1) &\le t_d(L_3-L). \end{aligned}$$

Using this inequality, we can eliminate x_{d2} because larger x_{d2} is smaller the average copy count is when all other parameters L_1, L_2, L_3, x_{d1} are kept constant, so using the above inequality as equation, we obtain:

$$x_{d2} = \frac{t_d(L_3-L)-x_{d1}(L_2-L_1)}{L_3-L_2}$$

Furthermore, the average copy count used in this case is:

$$\begin{aligned} c_3(L_1,L_2,L_3,x_{d1}) \approx\ & L_1+(L_2-L_1)e^{-\alpha L_1 x_{d1}} \\ & +(L_3-L_2)e^{-\alpha L_2(x_{d2}-x_{s1})} \end{aligned}$$

Continuing in the same way as in the two period case (i.e., substituting x_{s1}, x_{d2}, taking partial derivative and comparing to zero), we can obtain the formula for optimum x_{d1}.

$$x_{d1} = \frac{\alpha t_d L_2(L_3-L)+\log(L_1/L_3)(L_3-L_2)}{\alpha L_2(L_3-L_1)}$$

Then, we can easily obtain formula $c_3^*(L_1, L_2, L_3)$ by substituting x_{d1} with its optimum. Since $L_1 < L < L_3$ and $L_1 \le L_2 \le L_3$ and all these values are integers, by enumeration, we can simply find out the (L_1, L_2, L_3) combination that gives the minimum copy count for a given L.

IV. Simulation Results

In our simulations, we implemented the original Source Spray and Wait algorithm using a Java based visual simulator. We deployed 100 mobile nodes including the sink onto a torus of the size 300 m by 300 m. All nodes (except the sink that has high range for acknowledgment broadcast) are assumed to be identical and their transmission range is set at R = 10 m. Nodes move according to random direction mobility model [10] model. We have created messages at randomly selected source nodes for delivery to the sink node whose initial location is also decided randomly. Then, we collected some useful statistics from the network. The results are averaged over 1000 runs.

We have computed the appropriate combination of copy counts L_i for each period i analytically and tested it with our simulations. Table I shows the values of optimum L_i's for different L values.

We have calculated the average number of copies used by both simulations and the theory when this optimum L_i combination is used. Figure 5 and Figure 6 present these values for different L values, with two periods and three periods, respectively. In the two period case, results are very close to each other, however in the three period case, the difference gets bigger because in our analysis we ignored the effect of spraying phase. When number of periods increases, period lengths get smaller, so the effect of spraying phase on the cumulative distribution function increases.

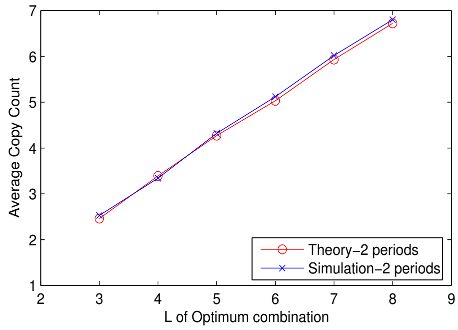

Fig. 5. The average copy count comparison for theory and simulation when the optimum value in 2 period case is used.

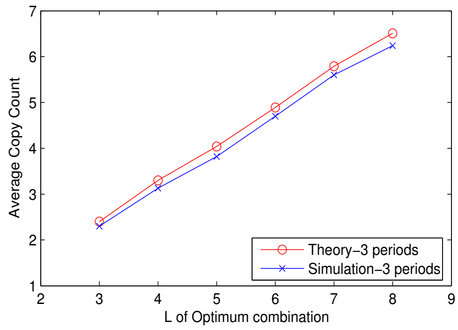

Fig. 6. The average copy count comparison for theory and simulation when the optimum value in 3 period case is used.

To compare the performance of our algorithm with the original spraying algorithm, we have measured some metrics of both of them via simulations. In these simulations, we used our algorithm with two periods. Figure 7 shows the the average value of delivery delay for messages. Figure 8 shows the average time of completing spraying. This value does not contain the average of cases when the message is delivered before spraying of all potential copies. In Figure 9, we show the success rate which is actually the percentage of all simulations that have delivery time less than or equal to the given deadline t_d.

When we look into these three graphs, we observe that our time based spraying algorithm incurs higher average delay but it achieves the same deliver rate before the deadline as the standard spraying algorithm. Moreover, since our scheme postpones the spraying of all copies to later times, it finishes spraying later than the standard Spray and Wait algorithm.

L	3	4	5	6	7	8
2 periods	(2,5)	(3,6)	(3,8)	(4,9)	(5,10)	(6,12)
3 periods	(2,3,6)	(2,4,7)	(3,5,9)	(4,6,10)	(5,7,11)	(5,8,14)

TABLE I

The optimum L_i combinations that achieve the minimum average copy count while preserving the delivery rate before deadline of the original algorithm.

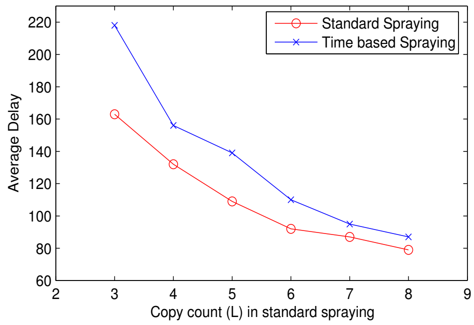

Fig. 7. The average delay comparison for standard spraying and time based spraying.

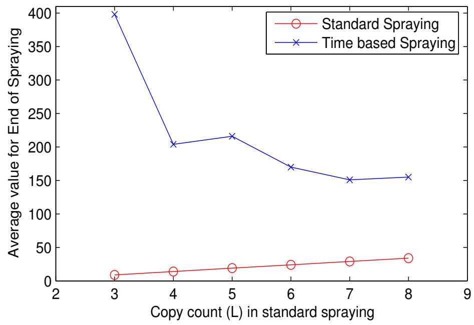

Fig. 8. A comparison of average times for end of spraying in standard spraying and time based spraying.

Finally, Figure 10 shows the improvement achieved by our algorithm in the average number of copies per message for different L values. While the two period case demonstrates about 16% benefit, the three period case shows higher improvement of about 20%.

V. Conclusion and Future Work

In this paper, we focus on the problem of routing for Delay Tolerant Networks in which the nodes are disconnected most of the time. We propose a time dependent spray and wait algorithm and evaluate its performance with simulations.

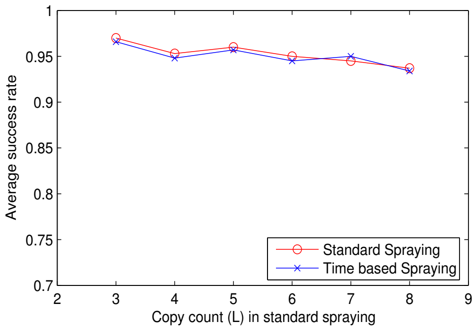

Fig. 9. The deliver rate comparison for standard spraying and time based spraying.

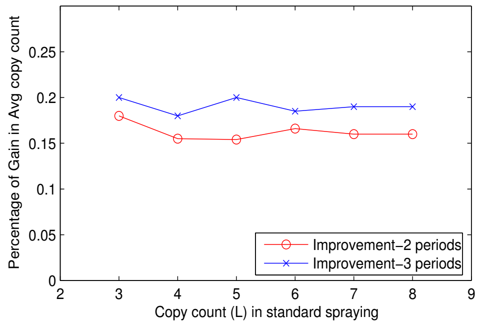

Fig. 10. The improvement obtained by the time based spraying in the number of average copies used.

We observed that the average number of copies used by our algorithm is lower than in the original spray and wait algorithm.

We applied our algorithm with just two and three period cases. In future work, we plan to apply it to binary spraying and consider cases with more periods. Furthermore, we also plan to apply our algorithm to a real test bed such as a disconnected bus network.

References

[1] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. S. Peh, and D. Rubenstein, *Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with zebranet*, in Proceedings of ACM ASPLOS, 2002.

[2] *Disruption tolerant networking*, http://www.darpa.mil/ato/solicit/DTN/.

[3] A. Vahdat and D. Becker, *Epidemic routing for partially connected ad hoc networks*, Duke University, Tech. Rep. CS-200006, 2000.

[4] J. Ott and D. Kutscher, *A disconnection-tolerant transport for drive-thru internet environments*, in Proceedings of IEEE INFOCOM, 2005.

[5] A. Lindgren, A. Doria, and O. Schelen, *Probabilistic routing in intermittently connected networks*, SIGMOBILE Mobile Computing and Communication Review, vol. 7, no. 3, 2003.

[6] 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.

[7] T. Small and Z. Haas, *Resource and peroformance tradeoffs in delat tolerant wireless networks*, in Proceedings of ACM SIGCOMM workshop on Delay Tolerant Networking (WDTN), 2005.

[8] T. Spyropoulos, K. Psounis,C. S. Raghavendra, *Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks*, ACM SIGCOMM Workshop, 2005.

[9] K. Harras, K. Almeroth, and E. Belding-Royer, *Delay Tolerant Mobile Networks (DTMNs): Controlled Flooding Schemes in Sparse Mobile Networks*, In IFIP Networking, Waterloo, Canada, May 2005.

[10] T. Spyropoulos, K. Psounis,C. S. Raghavendra, *Performance Analysis of Mobility-assisted Routing*, MobiHoc, 2006.

[11] C. Mascolo and M. Musolesi, *SCAR: Contextaware Adaptive Routing in Delay Tolerant Mobile Sensor Networks*, Proceedings of International conference on Wireless communications and mobile computing, 2006.

[12] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, *MaxProp: Routing for Vehicle-Based Disruption- Tolerant Networks*, In Proc. IEEE Infocom, April 2006.

[13] A. Balasubramanian, B. N. Levine, A. Venkataramani, *DTN routing as a resource allocation problem*, ACM SIGCOMM, 2007.

[14] T. Spyropoulos, K. Psounis,C. S. Raghavendra, *Efficient Routing in Intermittently Connected Mobile Networks: The Multiple-copy Case*, IEEE/ACM Transactions on Networking, 2008.