arXiv:cs.NI/0609068 v1 12 Sep 2006

The heterogeneity of inter-contact time distributions: its importance for routing in delay tolerant networks

Vania Conan*, Jérémie Leguay*+, Timur Friedman+
*Thales Communications
+Université Pierre et Marie Curie, Laboratoire LIP6–CNRS

September 12, 2006

ABSTRACT

Prior work on routing in delay tolerant networks (DTNs) has commonly made the assumption that each pair of nodes shares the same inter-contact time distribution as every other pair. The main argument in this paper is that researchers should also be looking at heterogeneous inter-contact time distributions. We demonstrate the presence of such heterogeneity in the often-used Dartmouth Wi-Fi data set. We show that the heavy-tailed distribution across all nodes, observed in previous work, can be explained as a composition of the distributions for each pair of nodes, and that these individual distributions are rarely heavy-tailed. We also show that DTN routing can benefit from knowing these distributions. We first introduce a new stochastic model focusing on the inter-contact time distributions between all pairs of nodes, which we validate on real mobility data. We then analytically derive the mean delivery time for a bundle of information traversing the network for simple single copy routing schemes. The purpose is to examine the theoretic impact of heterogeneous inter-contact time distributions. Finally, we show that we can exploit this user diversity to improve routing performance. Based on the analytic model, we define an improved single copy "Spray and Wait" scheme that we compare to other routing schemes on the same real mobility data.

1 INTRODUCTION

In delay tolerant networks (DTNs) [4], nodes are mobile and have wireless networking capabilities. They are able to communicate with each other only when they are within transmission range. The network is sparse and suffers from frequent connectivity disruptions, making the topology only intermittently and partially connected. This means that there is a very low probability that an end-to-end path exists between a given pair of nodes at a given time. End-to-end paths can exist temporarily, or may sometimes never exist, with only partial paths emerging. This paper addresses the extreme case, where only temporal paths exist. We call such networks *temporal DTNs*, or *t-DTNs*. When a node in a t-DTN receives a "bundle" of information from a neighboring node, there is a non-negligible interval before it contacts another node and has the opportunity to relay the bundle.

Prior work on routing in t-DTNs has commonly made the assumption that each pair of nodes shares the same inter-contact time distribution as every other pair. The main argument in this paper is that researchers should also be looking at cases in which inter-contact time distributions are heterogeneous.

We show, on the well known Dartmouth Wi-Fi data set [6], that despite the existence of a heavy-tailed distribution when inter-contact times are considered in the aggregate, a large portion of the node pairs present inter-contact time distributions that can be well fitted by an exponential distribution. We found these distributions to be heterogeneous, with a wide variation in exponents. Chaintreau et al. [2] posit that there might be heterogeneity, but we show it and characterize it. We also show how exponential distributions can be composed to yield the heavy-tailed distributions that Chaintreau et al. observed. As we shall see, the heterogeneity that we highlight allows us to usefully extend the work of Spyropoulos et al. [11, 10], which analyzes numerous routing schemes for t-DTNs, but that uses mobility models that yield homogeneous distributions.

We also provide the first formal analysis of the impact of heterogeneous exponential inter-contact time distributions on simple single-copy routing schemes. We show that routing strategies can benefit, in terms of delay, from this heterogeneity, and in particular from knowing these distributions. A node can choose among possible relay nodes based upon their expectations for meeting other relays or the destination.

2 INTER-CONTACT TIME MODEL

This section presents the model we use to analytically derive the delay expectations for the routing protocols we study.

2.1 Exponential t-DTNs

We consider a network composed of n nodes. Let's first look at the inter-contact time between two individual nodes (i, j): $t_{ij}^1 < t_{ij}^2 < t_{ij}^3 < ...$ are the successive instants at which a contact between i and j occurs.

$$\tau_{ij}^k = t_{ij}^{k+1} - t_{ij}^k \quad (1)$$

is the inter-contact time between the k^{th} and $(k+1)^{th}$ contact instants.

We assume that the τ_{ij}^k are samples from independent and identically distributed random variables that follow an exponential law with parameter λ_{ij}, which we note $\tau_{ij}^k = \tau_{ij} = \text{exponential}(\lambda_{ij})$. The mean inter-contact time between i and j is thus given by $E[\tau_{ij}] = 1/\lambda_{ij}$.

In the overall network, all n nodes are supposed to behave independently, so that the $n(n-1)/2$ pairwise inter-contact times τ_{ij} are independent exponential processes with different parameters. The τ_{ij} family of processes is symmetric and $\forall i, \tau_{ii} = 0$.

We can now proceed with the following definition: An exponential t-DTN of size n, n being the number of nodes in the t-DTN, is entirely and uniquely characterized by providing $n(n-1)/2$ strictly positive real parameters t_{ij} for all pair of nodes (i, j) with $i < j$. Parameter t_{ij} is the mean inter-contact time between node i and node j. By symmetry, $t_{ji} = t_{ij}$ for all nodes (i, j) with $i < j$. For all nodes i $t_{ii} = 0$. If two nodes l and k never meet, we use the convention $t_{kl} = \infty$. The inter-contact time processes are defined for all pairs of distinct nodes (i, j), $i \neq j$, by $\tau_{ij} = \text{exponential}(\lambda_{ij})$, with $\lambda_{ij} = 1/t_{ij}$.

2.2 Assumptions

The model abstracts away from all the spatial information that is essential in the analysis of mobile ad-hoc networks. There is no reference to geographic, localisation or any other such spatial information. There is also no reference to air interface parameters, quality of or contention on the links, etc. Node mobility is not explicit modelled: only its aggregated impact on the inter-contact time is taken into account.

The model focuses on the temporal dynamics of a DTN. In this way it provides a common framework to analyse very different DTNs. In particular it applies very well to social networks for which the position of nodes at a given time is not of primary importance. We believe this abstraction helps focus on the inherent characteristics of intermittent connectivity in DTNs.

The model makes a stationarity hypothesis with respect to node inter-contact time distributions. In other words, nodes behaviors are assumed to change on a slower scale than bundle exchanges.

We also suppose that nodes have infinite capacity. We do not study the impact of the load of the network, or of the limited bandwidth of the links, nor do we model the limitations due to buffer or queue overflows and corresponding bundle dropping strategies that nodes may require. In this respect, the results with the proposed model are upper bounds, but, as we will see, still provide valuable information and insight on how to route bundles in DTNs. We leave refinements of the model for future work.

3 Fitting the model

In the t-DTN model just elaborated, we assume that the inter-contact time distribution for each pair of nodes is exponential. The main reason is that it will allow us to go beyond asymptotic results and provide explicit formulas for the bundle delivery time, and other parameters, of different routing protocols. In this section we look at real data to evaluate how reasonable this hypothesis might be.

3.1 Exponential inter-contacts

To validate the hypothesis, we use real mobility traces inferred from the Wi-Fi access network of Dartmouth College [6]. The Wi-Fi scenario is not a perfect fit for DTN mobility. For instance, Wi-Fi nodes are typically turned off, transported, and then turned on again, thus missing potential contacts en route. However, the size, quality, and public availability of the data set make it nonetheless one of the best resources for this kind of study.

As we describe in prior work [9], we must select from the data, and make some assumptions, in order to constitute a useful DTN mobility data set. We take the subset of users who are present in the network every day between January 26th 2004 and March 11th 2004, a period during which we expect mobility patterns to be fairly stationery. This data set contains 834 users, or nodes. Then we assume that two nodes are in contact if they are present at the same time at the same access point (AP). Finally, we filter these data to remove the well known “ping pong” effect. Wireless nodes, even non-mobile, can oscillate at a high frequency between two APs. This leads to a large number of inter-contact times that are close to zero and thus biases the inter-contact time distribution by introducing an artificially high slope close to the origin. To counter this, we filter all the inter-contact times below 1,800 seconds. This threshold was used by Yoon et al. [12] for the same purpose. We use this data set for the remainder of this paper.

Fig. 1 shows the distribution of λ_{ij} for all 74,848 source-destination pairs for which this can be calculated. We see that the distributions are heterogeneous, with lambda varying over three orders of magnitude.

We test for whether the inter-contact process between any two nodes can be modelled by an exponential process with a parameter $\lambda = 1/\tau$, where τ is the mean inter-contact time. We use the Cramer-Smirnov-Von-Mises [3] hypothesis test. For each pair (i, j), we compare the cumulative distribution I_N^{ij} for the N inter-contacts observed and the hypothesis function whose cumulative distribution is $F_{ij}(x) = 1 - \exp(-\lambda_{ij}x)$. We also compare I_N^{ij}

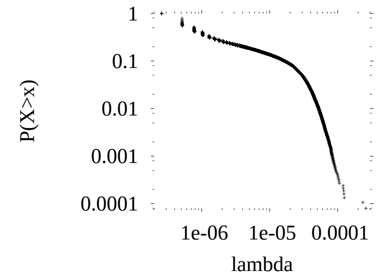

Figure 1: Distribution of λ_{ij}.

with that of a power law distribution. Note that we only perform the computation for pairs that show a sufficient level of connectivity by having a mean inter-contact time lower than one week and that have more than 20 contacts. We identify 8,402 pairs to be exponentially distributed and 28 with a power law which makes respectively, 62.3% and 0.2% of the 13,482 pairs tested.

It is clearly more reasonable, in this data set, to model pairwise inter-contact time distributions as exponential rather than power law. As we've examined only one data set, albeit an often-used one, we cannot draw many conclusions about what will be revealed elsewhere. It is reasonable to expect that other mobility traces in campus environments will show similar characteristics. However, it is surprising that a memoryless process seems to be at work in such a high proportion of node pairs in an environment in which one would expect some temporal correlations. We hope this will be a spur to study these distributions in other data sets.

3.2 Power laws

Chaintreau et al. [2] observed that aggregated inter-contact times follow power laws in a number of DTN mobility traces (also including one based on the Dartmouth data). Fig. 2 shows that, for our data set, the cumulative distribution of aggregated inter-contact times also follows a power law of the form $f(x) = ax^{\alpha}$, with exponent $\alpha = -0.16$ and scale parameter $a = 3.45$.

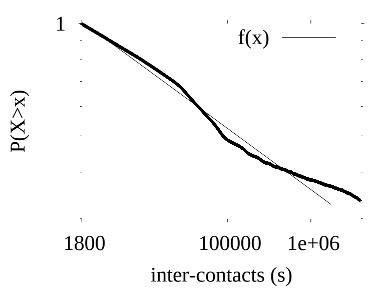

Figure 2: Distribution of inter-contacts.

Let's now consider what happens for pure exponential t-DTNs. Since all pairwise inter-contact time distributions are exponentially distributed, under which conditions do the aggregated inter-contact time distributions follow a power law, or is the pairwise exponential assumption too strong to yield a power law in the aggregate?

Let Θ be the aggregated inter-contact time for all pairs of nodes, and let $p(\lambda)$ be the probability distribution of the λ parameters:

$$P(\Theta > t) = \int_{\lambda=0}^{\infty} e^{-\lambda t} p(\lambda) d\lambda \qquad (2)$$

What eqn. 2 says is that, for exponential t-DTNs, the aggregated inter-contact time distribution is fully characterized by the distributions of the λ parameters, and thus of the t_{ij} matrix. More precisely, the tail cumulative distribution of the aggregated inter-contact times is given by the Laplace transform of the distribution p of the λ parameters.

A Pareto law of the form $(\frac{a}{t+a})^{\alpha}$, with shape parameter $\alpha > 0$ and scale parameter $a > 0$, is observed if and only if the λ follow a gamma distribution $p(\lambda) = \frac{\lambda^{\alpha-1} a^{\alpha} e^{-a\lambda}}{\Gamma(\alpha)}$.

To verify this on the data set we proceed in the following way: We estimate parameters α and a from the cumulative distribution of the λ parameters for pairs that were shown to follow an exponential behavior (the ones that pass the Cramer hypothesis test). We find $a = 113766.9$ and $\alpha = 2.26$. Fig. 3(a) shows the estimated cumulative gamma distribution $g(x)$ with the experimental lambda cumulative distribution. Then, we plot in Fig. 3(b) the corresponding power-law $h(x)$ with cumulative distribution of aggregated inter-contact times. As one can see, the two experimental curves roughly fit the theoretical curves.

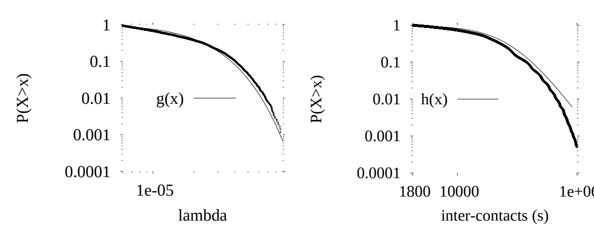

(a) λ_{ij} (b) Inter contacts

Figure 3: Distributions with exponential pairs.

What this result shows is that when one considers an exponential t-DTN, we can regain the power law behavior for the aggregated inter-contacts when the distribution of the parameters is a gamma, which is the case in the data we used.

4 SINGLE COPY ROUTING STRATEGIES

Having defined a stochastic model that is realistic for the data set under study, we now examine different simple single copy routing strategies. We derive analytical formulas that we will use to study the impact of heterogeneous λ_{ij} parameters on routing.

In all routing strategies, we consider that nodes know all pairwise mean inter-contact times for all nodes in the network, i.e., each node knows the λ_{ij} matrix. This knowledge could be diffused through an epidemic type of routing, or learned by each node from past contacts.

We consider that the contact time is negligible with respect to the inter-contact time (a reasonable hypothesis for several data sets [2]). For our purposes, then, contacts are instantaneous, and any bundles are transferred during that instant of contact.

4.1 Wait scheme

Under the Wait routing strategy (or 'direct communication'), the source node s waits until it meets d, the destination, to deliver the bundle in one hop.

If the bundle is injected at time t, its delivery time is equal to R^t_{sd}, the remaining inter-contact time before the next contact between nodes s and d. The memoryless nature of exponentials implies that R^t_{sd} also follows an exponential distribution with the same parameter. The mean expected delivery is thus given by:

$$E[D^w_{sd}] = 1/\lambda_{sd} \tag{3}$$

This straightforward result gives an upper bound on the delivery time that a routing strategy should meet, since the Wait strategy is the most rudimentary one hop single copy scheme.

4.2 MED

The Minimum Expected Delay (MED) routing strategy was first introduced by Jain et al. [7]. This strategy, similar to source routing, defines which path the bundle will follow from s to d, that is, the ordered list of intermediate relay nodes it will have to go through. The list is chosen to provide minimum expected end-to-end delay.

If a path is given by the following ordered list of nodes $r_0 = s < r_1 < r_2 < r_3 < ... < r_{n-1} < r_n = d$, and relaying occurs at time instants $t_1 < t_2 < ... < t_n$, the total delivery time along path $(s, r_1, r_2, ..., r_{n-1}, d)$ is given by the remaining inter-contact time after each relaying instant, that is:

$$D^{med}_{s,r_1,r_2,...,r_{n-1},d} = R^{t_1}_{sr_1} + R^{t_2}_{r_1r_2} + ... + R^{t_n}_{r_{n-1}d} \tag{4}$$

Using the fact that $E[R_{r_ir_j}] = 1/\lambda_{ij}$, the expected delivery time along the path is thus given by:

$$E[D^{med}_{s,r_1,r_2,...,r_{n-1},d}] = 1/\lambda_{sr_1} + 1/\lambda_{r_1r_2} + ...1/\lambda_{r_{n-1}d} \tag{5}$$

Finding the optimal path thus amounts to finding a lowest-weight path between nodes s and d in a graph in which the weight on each link (i, j) is defined as $1/\lambda_{ij}$. Dijkstra's algorithm can be used.

The weakness of the strategy, as already mentioned by Jain et al., is that it requires each node to wait for the next relay in the precomputed path. A node cannot take advantage of an opportunistic contact with a node that has a lower cost path than does the predesignated next hop node.

4.3 Spray and Wait routing

The Spray and Wait strategy was first introduced by Grossglauser and Tse [5], and is designed to take advantage of opportunistic contacts. It consists of two steps. First the source node uses the first nodes encountered as relays to the destination. This is the "spraying" step. A relay node then uses the "wait" strategy to relay the bundle, i.e. it waits until it meets the destination to deliver the bundle. Here, we study the case where only one relay is used, which we designate 1-SW.

Let us first consider the spraying step. The bundle is injected at source s at time instant t. The first node r it encounters may be any of the $n-1$ other nodes $d, r_1, r_2, ..., r_{n-2}$ and the time X it takes to meet this first node is the infinum of the inter-contact times with all other nodes:

$$X = inf(R^t_{sd}, R^t_{sr_1}, ..., R^t_{sr_{n-2}}) \tag{6}$$

Since all $R^t_{sr_i}$ are independent exponentials with parameters λ_{sr_i}, we have (see [1, p.328]):

- The random index r of the first node encountered is independent of the first encounter time X
- X is exponentially distributed, with parameter: $\Lambda_s = \lambda_d + \sum_{i=1}^{n-2}(\lambda_{sr_i})$
- $\Pr(First\ node\ encountered\ is\ r) = \frac{\lambda_{sr}}{\Lambda_s}$

This means that we can represent the spraying step as independently identifying the encountered node (with probability $\frac{\lambda_{sr}}{\Lambda_s}$) and adding an exponential waiting time with parameter Λ_s.

Two cases may arise: either the first node encountered r equals d, and s delivers the bundle, or $r \neq d$ and node r waits to meet node d to deliver the bundle.

The delivery time Z_d, when node d is encountered first is thus given by:

$$E[Z_d] = \frac{1}{\Lambda_s} \tag{7}$$

The delivery time Z_r along path r, i.e., conditioned on using node r as a relay, is thus the sum of the first encounter time X and the remaining delivery time between nodes r and d, and thus:

$$E[Z_r] = \frac{1}{\Lambda_s} + \frac{1}{\lambda_{rd}} \tag{8}$$

The total delivery time Z is computed by conditioning on all possible first encountered nodes $d, r_1, r_2, ..., r_{n-2}$, events whose probabilities are given by $\frac{\lambda_{sr}}{\Lambda_s}$:

$$E[Z] = \frac{\lambda_{sd}}{\Lambda_s} E[Z_d] + \sum_{i=1}^{n-2} \left(\frac{\lambda_{sr_i}}{\Lambda_s} E[Z_{r_i}]\right) \tag{9}$$

After simplification, we can state the following result: In a network composed of n nodes, the 1-SW routing algorithm delivers a bundle from source s to destination d with mean delivery time given by:

$$E[D_{sd}^{1-sw}] = \frac{(1 + \sum_{r \neq s, r \neq d} \frac{\lambda_{sr}}{\lambda_{rd}})}{\sum_{r \neq s} \lambda_{sr}} \tag{10}$$

5 Comparing routing protocols

In this section, we first show that all the routing strategies are equivalent when inter-contact time distributions are homogeneous among node pairs. Then, we show through simulations using the Dartmouth data that routing performance improves when the protocols take into account the heterogeneity of inter-contact time distributions.

5.1 Homogeneous case

In the homogeneous case, all nodes have the same mean inter-contact time $1/\lambda$: the $[\lambda_{ij}]$ matrix is thus given by $\forall i \neq j, \lambda_{ij} = \lambda$.

For a homogeneous exponential t-DTN of size n and parameter λ, we have:

1. for the Wait scheme (upper bound): $E[D^w] = 1/\lambda$
2. for the MED scheme: $E[D^{med}] = 1/\lambda$
3. for the 1-SW scheme: $E[D^{1-sw}] = 1/\lambda$

Result 1 follows from eqn. 3, concerning the Wait scheme. For MED, adding one relay to a path increases its delay by $1/\lambda$, so MED is identical to the Wait scheme (result 2). Result 3, for 1-SW, follows from eqn. 10 (see also Grossglauser and Tse [5]).

These results are not all new, but collectively they shed an interesting light on the homogeneous case. Generally speaking, they are rather negative for single copy schemes. The three single copy routing strategies that we studied give the same expected delivery time.

We conjecture that in a homogeneous exponential t-DTN no feasible single copy scheme does better than $1/\lambda$. For a given realisation of the joint random process of inter-contact times, the bundle follows a t-path between source and destination. Each node that possesses the bundle has to choose to which node to give it next. A feasible routing scheme must base its decision only on present or past information. However, an exponential process is memoryless, so past information is no help, and all nodes are seemingly equivalent as relays.

5.2 Heterogeneous case

This section looks at routing protocols that take into account heterogeneity in inter-contact time distributions. In this context, we present 1-SW*, a variation of 1-SW. Instead of spraying its bundle to the first node that it encounters, the source node s sprays only to nodes in a subset R. We call this a 1-SW^R scheme. Following the same line of reasoning as in Sec. 4.3, and defining $1/\lambda_{dd} = 0$, one finds that the expected delivery time is given by:

$$E[D_{sd}^{1-sw^R}] = \frac{(1 + \sum_{r \in R} \frac{\lambda_{sr}}{\lambda_{rd}})}{\sum_{r \in R} \lambda_{sr}} \tag{11}$$

We define a 1-SW* scheme to be a 1-SW^R scheme having a subset R that minimizes $E[D_{sd}^{1-sw^R}]$. We also study 1-MED, which uses source routing, as does MED, but which allows at most one relay on the path from source to destination.

We performed simulations with the mobility traces used in Sec. 3 to see how the algorithms studied analytically perform in the case of heterogeneous mobility. We simulate the following protocols: Wait, 1-SW, 1-SW*, MED and 1-MED. We slightly modified 1-SW, to better compare it with 1-SW*: a node i is defined as a relay only if $\lambda_{id} > 0$, i.e., if it has a chance of meeting the destination.

As computing the optimal set R for 1-SW* has a complexity in O(2^n) with a branch and bound algorithm [8], we used several heuristics to find solutions in a reasonable amount of time, at the expense of not always finding the optimal set. The basic idea of the branch and bound method is to build a tree that explores all the possible combinations of nodes that could form R and to discard complete parts of the tree based on a decision criterion. The heuristics we added are the following: we do not explore the tree further than a depth of 5, and if we do not find better solutions in a root branch of the tree for 10 sec. we skip to the next branch.

We choose at random 100 different source destination pairs (s,d). Once we add the nodes involved in the sets R for all the pairs selected, we complete the whole set of nodes with random nodes among the 835 present in the data set to reach 400 in total.

	delivery ratio (%)	**A delay** (days)	**M delay** (days)	**th. delay** (days)	**hopcount** (hops)
Wait	15.6 ±2.6	18.2 ±4.7	13.1 ±7.9	11.2 ±1.1	1.0 ±0.0
1-SW	93.4 ±2.2	20.9 ±2.1	18.8 ±4.2	9.5 ±0.9	1.9 ±0.0
1-SW*	93.6 ±1.7	18.2 ±1.0	15.6 ±0.8	2.5 ±0.2	1.9 ±0.1
MED	2.1 ±0.7	0.2 ±0.2	0.0 ±0.0	0.9 ±0.2	1.9 ±0.4
1-MED	15.4 ±2.4	9.0 ±2.9	2.3 ±1.5	4.2 ±0.6	1.7 ±0.1

Table 1: Simulation results with Dartmouth data.

Table 1 presents the simulation results averaged over 5 runs with the 90% confidence levels that are obtained using the Student *t* distribution. It presents, for each of the protocols, the average delivery ratio, the average delay ("A delay") and the median delay ("M delay") computed over the delivered bundles, the average theoretical delay over all the bundles generated, and the average hop count, also obtained on delivered bundles.

What we can first see is that Wait only delivers 15.6% of bundles because most of the source, destination pairs selected at random satisfy $\lambda_{sd} = 0$. For 1-SW and 1-SW*, we choose relays *i* such that $\lambda_{id} > 0$, however, as λ values were computed over the entire data set, a node *i* may meet the destination for the last time before having met the source for the first time. As a consequence, we still have respectively 6.6% and 6.4% of bundles that were not delivered for 1-SW and 1-SW*. Clearly, MED and 1-MED also suffer from this by achieving only 2.1% and 15.4% delivery ratios. Even with knowledge of average inter-contact times, deciding on the sequence of relays at the source is clearly a disadvantage. The high delivery ratios of 1-SW and 1-SW* might be due to their opportunistic natures, which the other algorithms do not share.

Selecting relays among the set R of nodes that minimizes (or nearly minimizes) eqn. 11 results in 1-SW* having a median delay of 15.6 days, as compared to 18.8 days for 1-SW.

To summarize these simulations results, we see that taking into account the heterogeneity of inter-contact time distributions is of great interest for the design of routing solutions for t-DTNs. Even relatively simple strategies such as 1-SW* perform relatively well. We thus expect that more elaborate schemes, in terms of number of copies distributed or in terms of the number of hops allowed, to achieve even better performance.

6 DISCUSSION AND CONCLUSION

We have seen in this work that, in a widely-used t-DTN mobility data set, distributions of inter-contact times are not heterogeneous. We saw through a formal analysis that, in homogeneous t-DTNs, practical routing is possible. However, we have argued that if inter-contact times follow an exponential distribution then routing cannot make practical use of inter-contact time information. On the other hand, in the heterogeneous case, a simple routing strategy, 1-SW*, adapted from the Spray and Wait scheme, is capable of using the diversity of inter-contact time distributions to improve routing performance, measured in terms of average delay.

Clearly, our work, based as it is upon one data set, will benefit from validation against others, such as the contact traces obtained in the Intel iMote based experiments [2]. Also, work needs to be done to examine why a memoryless model fits so many node pairs in an environment in which one would expect to find more temporal correlations.

What we show in this paper has also implications for DTN simulations. Chaintreau et al.'s findings call into question the use of some simple models, such as random-waypoint. But it provides nothing with which to replace those models. It is one thing to know that a model should produce a power law result. It is another to propose a simple model that produces that result. Our work suggests a way forward. Perhaps mobility could be modeled through a composition of sets of nodes, each moving according to the random-waypoint model, but with different parameters for each set. This will produce an exponential distribution of inter-contact times for the nodes within each set, while ensuring the heterogeneity of inter-contact time distributions across sets. Whether such an approach can produce the desired distributions between nodes in different sets, and overall, remains to be tested.

REFERENCES

[1] P. Bremaud. *Markov Chains, Gibbs Fields, Monte Carlo simulation, and queues*. Springer, 1999.

[2] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on the design of opportunistic forwarding algorithms. In *Proc. INFOCOM*, 2006.

[3] W. Eadie. *Statistical Methods in Experimental Physics*. Elsevier Science Ltd, 1971.

[4] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proc. SIGCOMM*, 2003.

[5] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. *Transactions on Networking*, 10(4):477–486, August 2002.

[6] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proc. MobiCom*, 2004.

[7] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proc. SIGCOMM*, 2004.

[8] A. Land and A. Doig. An automatic method of solving discrete programming problems. *Econometrika*, 28(3):497–520, 1960.

[9] J. Leguay, T. Friedman, and V. Conan. Evaluating mobility pattern space routing for DTNs. In *Proc. INFOCOM*, 2006.

[10] T. Spyropoulos, K. Psounis, and C. Raghavendra. Multi-copy routing in intermittently connected mobile networks. Technical report, USC, 2004.

[11] T. Spyropoulos, K. Psounis, and C. Raghavendra. Single-copy routing in intermittently connected mobile networks. In *Proc. IEEE SECON*, 2004.

[12] J. Yoon, B. Noble, M. Liu, and M. Kim. Building realistic mobility models from coarse-grained traces. In *Proc. MobiSys*, 2006.