

CREST: An Opportunistic Forwarding Protocol Based on Conditional Residual Time

Sunil Srinivasa

Department of Electrical Engineering
University of Notre Dame, Notre Dame, IN
ssriniv1@nd.edu

Sudha Krishnamurthy

Deutsche Telekom Laboratories
Berlin, Germany
krishnadots@gmail.com

Abstract—Opportunistic forwarding protocols take advantage of contact opportunities to route data in intermittently connected environments. In these environments, a fully connected path between the source and destination may not always exist and the contact schedules of all the nodes are not known in advance. Hence, one of the key challenges for a node is to make effective forwarding decisions using only a limited knowledge of the contact behavior of the nodes in the network. Based on an analysis of human mobility traces that we collected from our office environment, we introduce a new link metric, conditional residual time, that accurately estimates the time remaining for a pair of nodes to meet using only the local knowledge of their past contacts. We then propose a distributed protocol called CREST, that uses the conditional residual time to opportunistically forward messages between pairs of nodes. Experimental results show that CREST has a lower end-to-end delay compared to protocols that depend on future contact schedules and global knowledge of the contact behavior across the network. Furthermore, by disseminating only a few additional copies of the message, the delivery ratio of CREST improves significantly and is comparable to that of the flooding protocol.

I. INTRODUCTION

The proliferation of wireless standards such as IEEE 802.11, Bluetooth, ZigBee, and other low power radio-based technologies has made it viable to equip almost any device with communication capabilities. The portability of such communication-enabled devices allows them to be embedded in common mobile entities, such as vehicles [1], [2], humans [3], and animals [4], thereby making them ubiquitous. This is a powerful concept that can be used to gather information and communicate that between any two end points, even in regions where no permanent networking infrastructure is available, such as under-developed (e.g. [5]) and hard-to-access environments (e.g. underwater sensor networks [6], [7]). However, one of the challenges in routing data in such networks is that a fully connected path between a source and destination may not always exist due to different factors, such as the mobility of nodes, low node density, short radio range, or power-saving modes.

Network architectures and protocol designs that deal with routing data in intermittently connected environments is an emerging area of research that is referred to as delay-tolerant networking (DTN) and in some cases, as opportunistic networking. Unlike traditional routing protocols that regard disconnections as exceptions, protocols designed for

intermittently connected environments have to be inherently tolerant to delays/disruptions. Nodes need to follow a store-and-forward approach and must be able to make forwarding decisions on-the-fly. While opportunism offers great potential, one of the challenges to be addressed for data delivery, especially when nodes are mobile, is to determine exactly how to take advantage of the interconnection opportunities offered by intermediate relays. This requires an analytical characterization of realistic mobility traces for obtaining an in-depth understanding of the mobility model, and further developing practical forwarding protocols based on the model.

A. Background

Previous work in this area can be classified along two dimensions. The first dimension primarily studies the impact of mobility on forwarding, based on an analytical characterization of human [8], [9] and vehicular [1] mobility traces, and provides useful insights on the design and performance of opportunistic communication protocols. The second dimension focuses on the design of forwarding protocols for DTNs, but does not explicitly deal with the characterization of mobility traces. These protocols vary in the amount of information that nodes require to take advantage of the contact opportunities and in how they obtain that information. On the one hand, a node requires no knowledge and can either take advantage of every contact opportunity it has to forward the data (as in the flooding protocol [10]) or instead, can ignore every forwarding opportunity and deliver the data directly to the destination. On the other hand, a node may use complete future contact schedules to route data to the destination (e.g. MED [11]). In between these two extremes, several protocols have been designed that make informed decisions based on mobility-based metrics, such as the mean estimated expected delay (MEED [12]), delivery probability (PROPHET [13]), contact distributions [14], as well as metrics based on the social structure of the network [15], [16].

B. Overview and Contributions

The main contributions of this paper are fourfold. First, we collected mobility traces of people in our office to study the impact of human mobility on opportunistic forwarding. Like the campus and conference environments that have been used previously for a similar purpose [8], [9], office environments

too have a well-connected communication infrastructure. However, the motivation behind collecting these traces is the design of proximity-based communication protocols, where the contact between two mobile nodes or the communication devices that they carry presents an opportunity for data dissemination. Hence, the purpose of obtaining the traces is to model the intermittent connectivity arising from human mobility and use that as a basis for designing opportunistic protocols. Furthermore, since mobility in different environments may vary, the traces can be used to compare the impact of mobility on opportunistic forwarding in different environments.

Second, based on the mobility traces, we characterize the aggregate inter-contact duration (ICD) for our dataset in Section III. The ICD is a measure of the intermittent connectivity arising from the mobility of the nodes. We show that the aggregate ICD exhibits a composite behavior, which conforms to the findings reported in [9], [17]. We then analytically derive the average end-to-end delay for the direct hop and flooding schemes, which provide bounds on the delay performance of other opportunistic protocols. Third, we characterize the individual pairwise ICDs and observe that most of the pairs follow a log-normal distribution, but with different parameters.

Fourth, we introduce a new link metric, conditional residual time, that accurately estimates the remaining time for a pair of nodes to meet, based on their last time of contact. Using this metric as a basis for computing the delay in forwarding a message across a link, we propose CREST, a distributed opportunistic forwarding protocol that efficiently exploits the heterogeneous contact behavior in the network. Experimental evaluations in Section V using realistic mobility traces show that despite using only local knowledge of past contacts, CREST performs well and has lower end-to-end delay and higher delivery ratio compared to protocols like MEED [12] and MED [11] that make use of global knowledge. Moreover, CREST incurs minimal overhead compared to flooding [10] and MEED [12].

II. EXPERIMENTAL DATASET

In order to collect human mobility traces, we chose 52 participants belonging to different functional groups in our office: researchers from two different research groups, project leaders in the business division, department managers, system administrators, administrative staff, and student interns. The participants clipped on an Ekahau location-tracking tag [18] while they were in the office premises. These tags are small, less disruptive, and hence, more likely to be co-located with their owners than bulkier portable devices, like laptops. Therefore, they can capture most of the contact opportunities. The Ekahau real-time locationing system makes use of RSSI fingerprints from Wi-Fi access points within the office building to determine the current location of a tag. The location coordinate of a tag is specified by the floor that a node is on (participants were distributed across different floors), and its X and Y coordinates on that floor. For our study, we use the location traces collected over a one month period.

Previous work derives the contact information from the logs of Wi-Fi access points (e.g. [8]). One of the issues in using

this approach is that two devices that are attached to the same access point may be regarded as being in contact, even though the devices may themselves be out of range with respect to each other. In our work, we define two individuals to be in contact during a period of time, if the location coordinates reported by their respective tags indicate that they are located on the same floor and within 5 meters of each other. This value of the distance threshold is suitable for our open office environment, where people sit at a close distance. However, by simply choosing other appropriate threshold values, the same location traces can be used to model the contact behavior of different opportunistic transfer devices.

III. EXPERIMENTAL DATASET CHARACTERIZATION

In this section, we characterize the contact behavior in our office environment based on the mobility traces that we collected (see Section II). First, we analyze the aggregate ICD and determine bounds on the delay performance of DTN protocols. Second, we model the pairwise ICD and propose a new link metric that helps nodes to make effective forwarding decisions in DTNs. We now present the notation and assumptions that we use in our analysis.

Let the mobile network consist of N nodes, each associated with exactly one device that can be used for content transfer. We assume a slotted time model $t = 0, 1, 2, \dots$. For an arbitrary node pair $(i, j) \in \{1, \dots, N\} \times \{1, \dots, N\}$, define the contact process $C_{(i,j)}(t)$ as

$$C_{(i,j)}(t) = \begin{cases} 1 & \text{if } i \text{ and } j \text{ are in contact during slot } t \\ 0 & \text{otherwise.} \end{cases}$$

Let $0 < t_1 < t_2 < \dots < t_n < \dots$ denote the increasing sequence of time slots when $C_{(i,j)} = 1$. The inter-contact duration (ICD) is defined as the time elapsed between two successive contacts of a pair of nodes. Thus, for the pair (i, j) , the ICD after the k^{th} contact is simply $t_{k+1} - t_k$. For mathematical tractability, we assume that $\forall (i, j)$, $C_{(i,j)}$ is a renewal process, meaning that the ICDs associated with different node pairs are independent and identically distributed (i.i.d.). We denote the random variable (r.v.) representing the ICD for the pair (i, j) by $T_{(i,j)}$.

The residual ICD (or residual time) between devices i and j at time slot t , denoted by $r_{(i,j)}(t)$, is formally defined as

$$r_{(i,j)}(t) = \min \{t' - t : t' > t \text{ and } C_{(i,j)}(t') = 1\}.$$

It is a measure of the time remaining for the next contact between i and j , and is used to model the waiting time of a message in a node's buffer. We denote the r.v. representing the residual time for the pair (i, j) by $R_{(i,j)}$.

For the remainder of this paper, we assume that when two devices are in contact, any amount of information can be exchanged and that the buffer sizes of nodes are large enough that no messages are dropped. We also assume that the time taken to transfer data between devices in contact is negligible compared to the waiting time for the next contact opportunity, and that the latter has a more significant impact on the end-to-end delay.

A. Aggregate Inter-contact Durations

In this section, we characterize the complementary cumulative distribution function (CCDF) of the aggregate ICD, which is defined as the CCDF of the inter-contact durations aggregated over all device pairs over the measurement period. We then analytically study the delay performance of two well-known forwarding schemes: *direct-hop* transmission and *flooding*. For the remainder of this subsection, we assume that the ICD across all device pairs are i.i.d. and simplify the notation by dropping the subscript (i, j) .

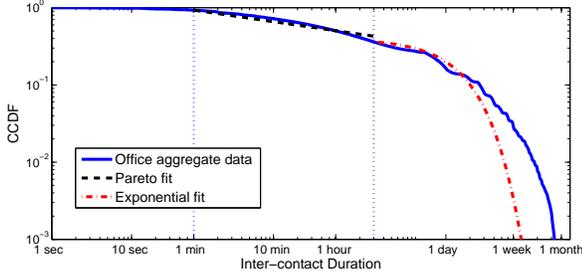


Fig. 1. CCDF of the aggregate ICD for the office dataset, and the corresponding Pareto and exponential fits. The two dotted vertical lines correspond to $t_m = 1$ minute and $t_c = 3$ hours.

Figure 1 shows the CCDF of the aggregate ICD for the office dataset. From the plot, we observe that the aggregate ICD exhibits a Pareto (or power-law) distribution up to a few hours, beyond which it decays exponentially. This shows that the behavior of the aggregate ICD in an office environment is similar to that observed in human mobility traces analyzed previously [9], [17]. In general, the CCDF of the aggregate ICD for this composite behavior can be written as

$$\bar{F}_T(t) = \begin{cases} 1 & 0 \leq t \leq t_m \\ (t/t_m)^{-\alpha} & t_m < t \leq t_c \\ (t_c/t_m)^{-\alpha} e^{-\lambda(t-t_c)} & t_c < t, \end{cases} \quad (1)$$

where t_m denotes the minimum value of T ; t_c is the characteristic (cut-off) time and denotes the time beyond which the CCDF decays rapidly; and α and λ are the parameters of the Pareto and exponential distributions, respectively. For the office dataset described in Section II, we determined that the values $\alpha = 0.1497$, $\lambda = 7.87 * 10^{-6}$, $t_m = 1$ min and $t_c = 3$ hours provide a reasonably good fit based on the K-S statistic.

We now analytically derive the average end-to-end delay for the direct-hop and flooding protocols, which provide upper and lower bounds, respectively, on the delay of other opportunistic routing schemes. We then evaluate the delays when the aggregate ICD exhibits a composite distribution, as specified by (1). Additionally, we provide some insights on the delay performance of these two protocols for the cases in which the distribution of the ICD is purely exponential or Pareto.

1) *Direct-Hop Transmission*: In this simple scheme, the source stores the message until it meets the destination directly and delivers the message in a single hop. Thus, the end-to-

end delay is simply the residual time between the source and destination nodes.

In order to derive the average end-to-end delay, we first note that when the contact process is a renewal process, the CDF of the ICD, $F_T(t)$, is related to the probability density function¹ of the residual ICD, $f_R(t)$, as follows [19, pp. 171-172]:

$$f_R(t) = \bar{F}_T(t)/\mathbb{E}[T], \quad (2)$$

where $\mathbb{E}[\cdot]$ denotes expectation (mean value). Accordingly, the n^{th} moment ($n = 1, 2, \dots$) of the residual time R is given by

$$\mathbb{E}[R^n] = \int_0^\infty t^n f_R(t) dt = \frac{1}{\mathbb{E}[T]} \int_0^\infty t^n \bar{F}_T(t) dt. \quad (3)$$

Thus, the average end-to-end delay for the direct-hop scheme is $\mathbb{E}[D_{\text{dh}}] = \mathbb{E}[R] = \int_0^\infty t \bar{F}_T(t)/\mathbb{E}[T] dt$.

When the CCDF of the aggregate ICD is specified by (1), we find that

$$\mathbb{E}[D_{\text{dh}}] = \frac{\frac{\alpha t_c^2}{2-\alpha} \left(1 - \left(\frac{t_m}{t_c}\right)^{2-\alpha}\right) + \frac{\lambda^2 t_c^2 + 2\lambda t_c + 2}{\lambda^2}}{\frac{2\alpha t_c}{1-\alpha} \left(1 - \left(\frac{t_m}{t_c}\right)^{1-\alpha}\right) + \frac{2\lambda t_c + 2}{\lambda}}. \quad (4)$$

Note that the exponential tail of the CCDF of the aggregate ICD guarantees a finite end-to-end delay value, irrespective of the other system parameters. Also, since (4) is an upper bound, the end-to-end delay of other DTN routing protocols is also finite. For the office dataset, the mean end-to-end delay using direct-hop transmission is computed from (4) to be about 35 hours.

Remarks:

- By letting $t_c \rightarrow \infty$ in (1), we can study the case where the aggregate ICD purely follows a Pareto distribution with parameter α . It is then straightforward to note that the end-to-end delay for the direct-hop scheme is infinite when $\alpha \leq 2$ and bounded otherwise.
- Setting $t_m = t_c \rightarrow 0$ in (1), the aggregate ICD takes the form of an exponential distribution with parameter λ . For this case, the end-to-end delay for the direct-hop scheme is equal to the mean of the exponential distribution, i.e., $1/\lambda$.

2) *Flooding*: In the flooding strategy [10], each node that has a copy of the message forwards a copy to every other node K that it meets, provided that K does not already have a copy. Since messages are flooded through every possible path from the source to the destination, the flooding strategy delivers messages with the shortest possible delay.

We now derive a closed-form expression for the end-to-end delay of the flooding protocol. For our analysis, we use the following notation: given a r.v. X , let $M_n(X)$ denote the r.v. $\min\{X_1, \dots, X_n\}$, where X_1, \dots, X_n are n i.i.d. random variables with the same distribution as X .

Let at some arbitrary time t , i distinct nodes (including the source node, but not the destination node) have a copy of the

¹Even though time is slotted, we can take the slot duration to be arbitrarily small, and treat T and R as continuous random variables.

message. We determine the additional time it takes for any of the remaining nodes to obtain a copy of the message. For each node having the message, the time elapsed before it meets any of the other $N - i$ nodes is the minimum of $N - i$ residual times (that are i.i.d.) and is represented by the r.v. $M_{N-i}(R)$. Therefore, the time elapsed before any of the i nodes meet any of the other $N - i$ nodes, i.e., the additional time it takes before $i + 1$ distinct nodes obtain copies of the message (given that i nodes already have the packet) is modeled by the r.v. $M_i(M_{N-i}(R))$. Let $d_i = \mathbb{E}[M_i(M_{N-i}(R))]$.

Now, note that the message is initially present only at the source node i.e., $i = 1$ at $t = 0$. Also, the destination node has a uniform probability of $1/(N - 1)$ of occurring at each of the $N - 1$ time instants at which a new node gets the message. Thus, the average end-to-end delay for the flooding protocol is equal to

$$\mathbb{E}[D_f] = \frac{d_1 + (d_1 + d_2) + \dots + (d_1 + d_2 + \dots + d_{N-1})}{N - 1},$$

which is equivalent to

$$\mathbb{E}[D_f] = \frac{1}{N - 1} \sum_{i=1}^{N-1} (N - i) \mathbb{E}[M_i(M_{N-i}(R))]. \quad (5)$$

We now show the procedure to evaluate (5) using only the CCDF of the residual time, \bar{F}_R . Indeed, by definition,

$$\begin{aligned} \bar{F}_{M_{N-i}(R)}(t) &= \Pr(R_1 > t, R_2 > t, \dots, R_{N-i} > t) \\ &= (\Pr(R > t))^{N-i} = (\bar{F}_R(t))^{N-i}, \end{aligned}$$

and consequently,

$$\bar{F}_{M_i(M_{N-i}(R))}(t) = (\bar{F}_R(t))^{(N-i)i}. \quad (6)$$

Putting together (5) and (6), the average end-to-end delay for the flooding protocol is obtained as

$$\mathbb{E}[D_f] = \frac{1}{N - 1} \sum_{i=1}^{N-1} \left[(N - i) \int_{\infty}^0 t \frac{d}{dt} (\bar{F}_R(t))^{N-i} dt \right]. \quad (7)$$

When the CCDF of the aggregate ICD is given by the composite distribution in (1), we use (2) to obtain,

$$F_R(t) = \frac{1}{\mathbb{E}[T]} \begin{cases} t & 0 < t \leq t_m \\ \frac{t^{1-\alpha} - \alpha t_m^{1-\alpha}}{(1-\alpha)t_m^\alpha} & t_m < t \leq t_c \\ \mathbb{E}[T] - \left(\frac{t_c}{t_m}\right)^{-\alpha} \frac{e^{-\lambda(t-t_c)}}{\lambda} & t_c < t, \end{cases}$$

Using the above expression in (7), the average delay for the flooding protocol is numerically calculated as 4.66 hours. The flooding protocol delivers messages about 7.5 times faster than direct-hop transmission.

Remarks:

- When the aggregate ICD follows a Pareto distribution with parameter α , the residual time is also Pareto-distributed, but with parameter $\alpha - 1$ [9]. Using this in (7), we obtain the mean end-to-end delay as

$$\mathbb{E}[D_f] = \frac{t_m}{N - 1} \sum_{i=1}^{N-1} \frac{(\alpha - 1)(N - i)^2 i}{\alpha^{(N-i)i} ((\alpha - 1)(N - i)i - 1)}.$$

The average end-to-end delay is infinite, if the Pareto parameter is ≤ 1 , i.e., when $(\alpha - 1)(N - i)i \leq 1$, for any $1 \leq i \leq N - 1$. Noting that the maximum value of $i(N - i)$ occurs at $i = N/2$ for even i , and at $i = (N - 1)/2$ for odd i , the end-to-end delay is infinite only when

$$\alpha \leq \begin{cases} 1 + 4/N^2 & \text{for even } N \\ 1 + 4/(N^2 - 1) & \text{for odd } N. \end{cases}$$

- When the aggregate ICD, T , is exponentially distributed with parameter λ , the residual time, R , also follows an exponential distribution with parameter λ . Thus,

$$\mathbb{E}[D_f] = \frac{1}{\lambda(N - 1)} \sum_{i=1}^{N-1} \frac{N - i}{i(N - i)} = \frac{1}{\lambda(N - 1)} \sum_{i=1}^{N-1} \frac{1}{i}.$$

From the above expression, we see that even in a network with just 10 nodes, flooding delivers messages about 3.2 times faster than direct-hop transmission. The end-to-end delay reduces further with increasing N , but at the cost of higher message overhead.

B. Pairwise Inter-contact Durations

The use of aggregate ICD helps in analytically evaluating the delay performance of some simple forwarding schemes and thus provides useful insights on the impact of using other protocols. However, in general, the aggregate ICD is not representative of the contact behavior of different pairs of nodes. For instance, device pairs that are in contact frequently have a lower average ICD than those that meet rarely. We remark that the individual pairwise ICDs provide a better idea of the contact behavior in the network. We now characterize the pairwise ICDs in human mobility traces and study their impact on opportunistic forwarding.

We used a distribution fitting software, EasyFit [20], to determine the best fit for the mobility traces that we collected. After fitting the pairwise ICDs to several types of distributions, including exponential, log-normal and Pareto, we observed based on the K-S statistic that the ICD between most pairs is fit best by the log-normal distribution, but with different parameters. Based on these results, we consider that $\forall(i, j)$, the pairwise ICD, $T_{(i,j)}$, is log-normally distributed with parameters μ_{ij} and σ_{ij} . Accordingly, the CCDF of the ICD for the device pair (i, j) is given by [21]:

$$\bar{F}_{T_{(i,j)}}(t) = \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left[\frac{\ln(t) - \mu_{ij}}{\sigma_{ij} \sqrt{2}} \right], \quad t \geq 0,$$

where $\operatorname{erf}(\cdot)$ is the error function. For our dataset, the distributions of the pairwise ICDs are heterogeneous with their means spanning over three orders of magnitude.

The characterization of contact behavior between individual pairs as log-normal distributions indicates that the pairwise ICDs are not memoryless. The above property implies that the residual time between a pair of nodes is dependent on their previous time of contact. Hence, the delay involved in forwarding a message between a node pair can be accurately

estimated by considering their last time of contact. Following this, we introduce the notion of conditional residual time (CRT), and use it to model the delay in forwarding messages across links in intermittently connected networks.

Conditional residual time: We define CRT as the time remaining before devices i and j meet, conditioned on the information that they last met t_{ij} time slots ago. Let $\hat{R}_{(i,j)}$ denote the r.v. representing the CRT between the device pair (i, j) . Formally, the CCDF of $\hat{R}_{(i,j)}$ is written as

$$\begin{aligned} \bar{F}_{\hat{R}_{(i,j)}}(t) &= \Pr\left(\hat{R}_{(i,j)} > t \mid T_{(i,j)} > t_{ij}\right) \\ &= \Pr\left(T_{(i,j)} > (t + t_{ij}) \mid T_{(i,j)} > t_{ij}\right) \\ &\stackrel{(a)}{=} \frac{\Pr\left(T_{(i,j)} > (t + t_{ij})\right)}{\Pr\left(T_{(i,j)} > (t_{ij})\right)}, \end{aligned} \quad (8)$$

where (a) is from the definition of conditional probability. When $T_{(i,j)}$ follows a log-normal distribution (as in the office dataset) with parameters μ_{ij} and σ_{ij} , we have

$$\bar{F}_{\hat{R}_{(i,j)}}(t) = \frac{1 - \operatorname{erf}\left(\frac{\ln(t_{ij} + t) - \mu_{ij}}{\sigma_{ij}\sqrt{2}}\right)}{1 - \operatorname{erf}\left(\frac{\ln t_{ij} - \mu_{ij}}{\sigma_{ij}\sqrt{2}}\right)}. \quad (9)$$

We now use the characterization of the CRT to propose a new metric that accurately models the delay in forwarding a message across a link in the network. This metric can then be used to represent the cost of a link and to efficiently forward messages from the source to the destination.

Indeed, a simple choice for the link metric is the mean or the average CRT. From (9), we obtain the mean CRT when $T_{(i,j)}$ is log-normal as:

$$\mathbb{E}[\hat{R}_{(i,j)}] = \frac{\exp\left(\mu_{ij} + \frac{\sigma_{ij}^2}{2}\right) \left(1 - \operatorname{erf}\left(\frac{\ln t_{ij} - \mu_{ij} - \sigma_{ij}^2}{\sigma_{ij}\sqrt{2}}\right)\right)}{1 - \operatorname{erf}\left(\frac{\ln t_{ij} - \mu_{ij}}{\sigma_{ij}\sqrt{2}}\right)} - t_{ij}. \quad (10)$$

The main shortcoming of using (10) is that the mean CRT is sensitive to extreme values of the data, particularly when the data size is small. For such cases, using the mean CRT does not accurately represent the cost of a link.

In order to circumvent this issue, we consider the *median* CRT, which is less sensitive to extreme values of the CRT, and therefore is a more robust link metric than the mean CRT. Let \tilde{t}_{ij} denote the median CRT, which is computed as $\tilde{t}_{ij} = \bar{F}_{\hat{R}_{(i,j)}}^{-1}(0.5)$. When the ICD is log-normally distributed, we have

$$\tilde{t}_{ij} = \exp\left(\operatorname{erf}^{-1}\left(\frac{1}{2} + \frac{1}{2}\operatorname{erf}\left(\frac{\ln t_{ij} - \mu_{ij}}{\sigma_{ij}\sqrt{2}}\right)\right)\right) \sigma_{ij}\sqrt{2} + \mu_{ij} - t_{ij}. \quad (11)$$

Figure 2 plots \tilde{t}_{ij} versus t_{ij} for the range of μ_{ij} and σ_{ij} values obtained from the characterization of the individual pairwise ICDs for the office dataset. For most pairs in our dataset, the value of μ_{ij} lies between 8.0 to 11.0, while σ_{ij} lies between 2.5 and 3.5. We observe from Figure 2 that when the ICD follows a log-normal distribution, the median CRT monotonically increases with the time elapsed since last

contact. Therefore, the longer it has been since the node pair (i, j) last met, the higher is the median time remaining for their next contact. However, note that the behavior of the median CRT may change depending on the distribution of the ICD. For example, when T_{ij} is exponential, the contact process is memoryless, and therefore \tilde{t}_{ij} is independent of t_{ij} . On the other hand, when the contact behavior between a pair of nodes is periodic, i.e., when $T_{(i,j)}$ is a constant, the residual time decreases with the elapsed time since their last contact. This can also be seen by letting $\sigma_{ij} \rightarrow 0$ in (11), so that $\tilde{t}_{ij} \rightarrow \exp(\mu_{ij}) - t_{ij}$, which decreases linearly with t_{ij} .

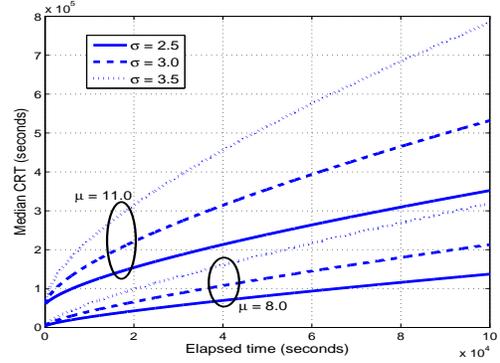


Fig. 2. \tilde{t}_{ij} versus t_{ij} for different values of the parameters μ_{ij} and σ_{ij} . When the pairwise ICD follows a log-normal distribution, the median CRT monotonically increases with time elapsed since last contact.

IV. CREST FORWARDING PROTOCOL

Having introduced a new metric to model the delay in forwarding messages across links in an opportunistic network, we now explain how it can be used by the nodes to make forwarding decisions. We propose CREST, an opportunistic protocol which performs forwarding decisions in a fully distributed manner, using only the local knowledge of the nodes that are in contact.

Algorithm 1 : CREST Forwarding Algorithm

- 1: **Inputs:** Source S , Destination D
 - 2: **Initialize:** Delivered = 0, ForwardingNode = $\{S\}$
 - 3: **while** Delivered == 0 **do**
 - 4: EncounteredNodes = set of all nodes that are currently in contact with the ForwardingNode;
 - 5: **if** $D \in$ EncounteredNodes **then**
 - 6: Forward message to D ; Delivered = 1;
 - 7: **else**
 - 8: PossibleRelays = {ForwardingNode} \cup {EncounteredNodes};
 - 9: **for all** nodes $i \in$ PossibleRelays **do**
 - 10: Compute the link metric, \tilde{t}_{iD} ;
 - 11: Send \tilde{t}_{iD} to ForwardingNode;
 - 12: **end for**
 - 13: NextHopNode = $\{k \mid \tilde{t}_{kD} \leq \tilde{t}_{iD}, \forall i, k \in$ PossibleRelays $\}$
 - 14: **if** NextHopNode \neq ForwardingNode **then**
 - 15: Forward message to NextHopNode;
 - 16: ForwardingNode = NextHopNode;
 - 17: **end if**
 - 18: **end if**
 - 19: **end while**
-

Algorithm 1 presents the steps involved in the CREST protocol. The goal of this algorithm is to make use of contact opportunities to forward a message from the source to the destination through one or more hops, while keeping the end-to-end delay low. Accordingly, when two or more nodes are in contact, the node having the lowest value of the median CRT with respect to the destination obtains the message. We now describe the CREST protocol in detail.

We refer to the node currently bearing the message as the forwarding node. When the forwarding node meets one or more nodes at a given time, each of the nodes that is in contact locally computes its median CRT with respect to the destination D using Equation (11). This involves two steps. First, each node i that is in contact, updates the values of μ_{iD} and σ_{iD} based on the knowledge of its past ICDs with D . Given the values of past ICDs for the node pair (i, D) as τ_1, \dots, τ_n , the log-normal parameters μ_{iD} and σ_{iD} are calculated as [21]:

$$\mu_{iD} = \ln \left(\mathbb{E}[\tau] - \frac{1}{2} \ln \left(1 + \frac{\sigma_\tau^2}{(\mathbb{E}[\tau])^2} \right) \right)$$

and

$$\sigma_{iD} = \sqrt{\ln \left(1 + \frac{\sigma_\tau^2}{(\mathbb{E}[\tau])^2} \right)},$$

where $\mathbb{E}[\tau] = \sum_{i=1}^n \tau_i / n$ and $\sigma_\tau^2 = \sum_{i=1}^n (\tau_i - \mathbb{E}[\tau])^2 / n$, are the arithmetic mean and variance of the ICDs. Second, each node i uses its knowledge of the time elapsed since its last contact with the destination, t_{iD} , and the values of μ_{iD} and σ_{iD} to compute its median CRT metric, \tilde{t}_{iD} (Line 10). The forwarding node then receives the computed link metric values from each of the nodes that it is currently in contact with (Line 11). If the forwarding node has the lowest value of the metric, no forwarding is done and the message is stored until its next contact opportunity. Otherwise, the message is forwarded to the node having the lowest median CRT (Line 15), which then becomes the new forwarding node. This process is repeated until the message is delivered to the destination.

In computing the link metric, \tilde{t}_{iD} , in the above description of CREST, we assume that the pairwise ICDs follow a log-normal distribution, which is the case for the mobility traces that we collected. However, the nodes can also use the steps presented in Algorithm 1 in other cases, by computing \tilde{t}_{iD} according to the distribution of the pairwise ICDs (using (8)).

The advantage of the CREST algorithm is that it enables the nodes to make forwarding decisions in a completely decentralized manner. Each node computes its median CRT with respect to the destination locally, using only the knowledge of its prior contacts with the destination. Moreover, CREST is a low overhead protocol and does not require the link costs to be disseminated, each time they are updated.

V. PERFORMANCE EVALUATION

In this section, we experimentally evaluate the performance of the CREST protocol (based on the median CRT metric) using simulations, and compare it with the following protocols: flooding [10], PROPHET [13], MEED [12], MED [11], and

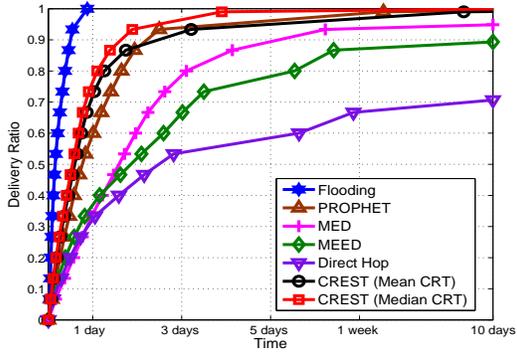
direct hop. The evaluation is based on the delivery ratio, end-to-end delay, message overhead (which is the total number of copies of a message transmitted in the network before it is delivered), and hop count (which is a measure of the path length between the source and destination). We first evaluate the protocols for the single copy case, where there is only one copy of every distinct message in the network at any given time. We then consider the case where the source forwards multiple copies of a message. We also evaluate how effective the median CRT is compared to the mean CRT metric in making forwarding decisions, based on an implementation of CREST that computes the link metric using the mean CRT (see (10)) instead of the median CRT.

Our evaluation is based on realistic human mobility traces collected from office (see Section II) and conference [3] environments. In sections V-A and V-B, we use the contact traces we collected over a period of four weeks in our office. During the simulation of each protocol, we generate 10000 messages in the network. The source and destination nodes for each of the messages are randomly chosen from the 52 participants in our dataset. Each message is generated at a randomly chosen time during the second of the four weeks. In V-C, we evaluate CREST using the human mobility traces collected from a conference environment, as part of the Hagggle project [3].

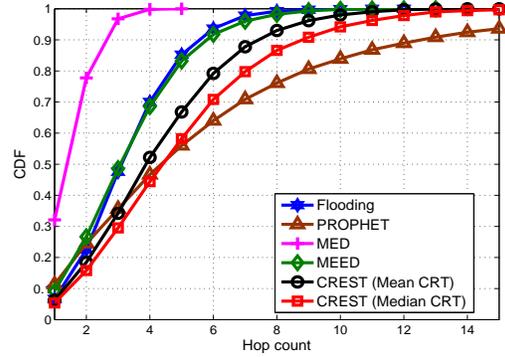
A. Single Copy

Figure 3(a) plots the CDF of the end-to-end delays for the different protocols, which also depicts their delivery ratios at different instants of time, since the time at which the message was generated. Figure 3(b) shows the CDF of the hop count for the delivered messages. The plots in Figure 3(a) show that the CREST protocol based on the median CRT has a higher delivery ratio and lower end-to-end delay compared to the protocol based on the mean CRT metric, although the latter has a slightly lower hop count. We infer that the median CRT is a more effective forwarding metric than the mean CRT.

We now compare the CREST protocol based on the median CRT with the remaining protocols. Evidently, the direct hop and flooding strategies provide the upper and lower bounds for the end-to-end delay and the delivery ratio. Direct hop is able to deliver only 35% of the messages in a day and about 50% in 3 days. Flooding delivers 100% of the messages within just 21 hours, whereas the CREST protocol (based on the median CRT) delivers almost 95% of the messages in about 2 days. However, note that the implementation of flooding, like the other schemes, does not impose any constraints on the buffer length. Additionally, flooding has a high overhead and on an average, results in 196 message transmissions. In contrast, the CREST protocol has a low message overhead of 5.3 transmissions of a message, on an average. While neither the flooding nor the direct hop scheme uses any knowledge about the network, we now compare the performance of CREST with protocols that use some knowledge of the network for making forwarding decisions.



(a) CDF of the end-to-end delay.



(b) CDF of the hop count.

Fig. 3. Comparison of opportunistic forwarding protocols for the office dataset: single copy. The delivery ratio along the y-axis is computed as the CDF of the end-to-end delay.

Both the MED [11] and MEED [12] protocols use the average residual time to model the cost of a link and construct a path with the shortest cost to forward a message from the source to the destination. However, while MED assumes that the future contact schedules between pairs of nodes are known in advance and performs source routing, MEED uses past contact information and makes forwarding decisions upon contact. Figure 3(a) shows that MED delivers only 34% of the messages in 1 day and 90% of the messages within 5 days. The CREST protocol, which uses only local knowledge, performs better than MED, even though MED uses full knowledge of future contact schedules. We explain this behavior by noting that unlike MED, where the complete path is decided at the source, in the case of CREST, the nodes update their link metrics upon an encounter and hence, the forwarding decisions are more adaptive to recent contact information. Also, contact behavior in intermittently connected environments may not be memoryless (see Section III-B). Hence, it is advantageous to take into account the immediate past contact time between a pair of nodes to estimate the time remaining until their next meeting. The median CRT completely captures this information, and hence is a more effective metric than the average residual time. We also observe from Figure 3(a) that CREST has a smaller end-to-end delay and a higher delivery ratio compared to MEED, which is able to deliver only 36% of the messages in 1 day and 78% of the messages in 5 days. Moreover, unlike MEED, CREST avoids the overhead incurred in the epidemic propagation of the link costs each time they are updated. Based on these results, we infer that compared to the average residual time metric used by MED and MEED, the median CRT estimates the delay in forwarding a message across a link more accurately and thus, enables the nodes in an opportunistic network to make more effective forwarding decisions.

We now compare the performance of CREST and the PROPHET protocol [13], which uses a history of encounters and transitivity to predict the delivery probability. Figures 3(a) and 3(b) show that the CREST protocol performs better than PROPHET, even though CREST does not incur the

overhead of exchanging routing tables. Within a one-day period, PROPHET delivers only 60% of the messages, while CREST delivers almost 80% of the messages. Moreover, the average hop count for PROPHET is 6.5, which is a little higher than the value of 5.5 for CREST. In PROPHET, the delivery probability of a node is aged with time and updated upon next contact, independent of the mobility model. For ICDs that follow distributions, such as log-normal or Pareto, the longer it has been since two nodes have met, the less probable it is that they will meet again soon (based on (8)). Hence, PROPHET performs reasonably well in these cases, although the aging parameters need to be chosen appropriately. However, PROPHET is not adaptive to different mobility models. For instance, if the contact behavior between a pair of nodes were to be periodic, then the longer it has been since their last contact, the smaller will be the residual time. In such a case, aging the delivery probability is not appropriate. In contrast, CREST does not explicitly age the delivery probability, but instead uses the characterization of the pairwise ICD to accurately model the delays across links.

B. Multiple Copies

With the exception of flooding, in each of the protocols considered in Figure 3(a), a node does not retain a copy of the message upon forwarding. However, since there is no guarantee that a pair of nodes in an intermittently connected environment will meet again, routing protocols designed for DTNs can improve the delivery ratio by replicating messages. Figure 3(a) shows that disseminating multiple copies of a message at each hop helps the flooding scheme to achieve a shorter delay and a higher delivery ratio. Similarly, the RAPID protocol [22] performs well by replicating packets at each transfer opportunity, but instead of flooding, RAPID models DTN routing as a utility-driven resource allocation problem and routes packets by replicating them in decreasing order of their marginal utility. The question we want to answer in this section is whether the CREST protocol can do as well as flooding, but with a small number of message copies, and if so, what is a good value for the degree of replication. Unlike

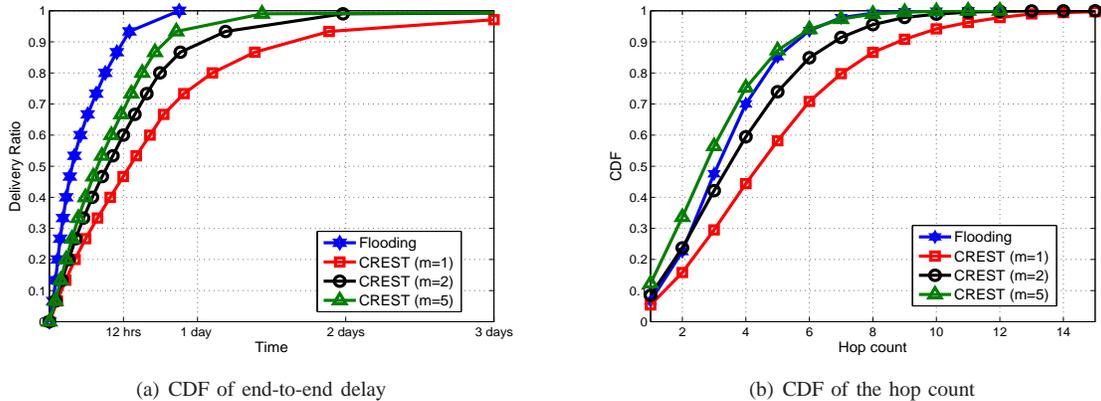


Fig. 4. Comparison of opportunistic forwarding protocols for the office dataset: multiple copies.

flooding and RAPID, however, we consider that the packets are replicated only at the source.

To answer the aforementioned question, we modify the CREST protocol, so that the source generates m copies of the message and forwards them to the first m nodes that it meets, which have a lower median CRT value to the destination than itself, provided that they do not already have a copy of the message. Note that the intermediate forwarding nodes do not replicate the message. Figure 4 compares the performance of the flooding scheme and the CREST protocol for $m = 1, 2$, and 5. The plots show that there is a steady improvement in the performance when m increases from 1 to 5. The time it takes to achieve 90% delivery ratio when $m = 1$ is about 40 hours. This drops to 24 hours when $m = 2$, and further to only 18 hours when $m = 5$. We did not observe significant improvement in the performance beyond $m = 5$. While flooding delivers 100% of the messages within 21 hours, the CREST protocol with $m = 5$ delivers as much as 95% of the messages in the same time, but at a much lower overhead. On an average, the flooding strategy disseminates 196 copies of a message through the network before it is delivered, while the corresponding values for the CREST protocol are only 5.34, 6.53, and 8.92, for $m = 1, 2$, and 5, respectively. Increasing the number of message copies from 1 to 5 also lowers the average hop count from 5.34 hops for $m = 1$, to 4.30 hops for $m = 2$ and further to just 3.46 hops for $m = 5$, which is lower than the average hop count value of 3.77 for the flooding strategy. Thus, the CREST protocol with only a few message copies is able to achieve end-to-end delays and delivery ratio comparable to that of flooding, while incurring a much lower overhead.

C. Performance Evaluation for a Conference Dataset

We also evaluate the performance of CREST using the mobility traces collected in a conference environment [3]. The conference dataset consists of the contact data logged by iMote devices carried by 41 conference attendees over a four day period. Each iMote device records the sightings of other iMote devices (termed internal contacts) as well as other types of Bluetooth devices (termed external contacts). For our evaluation, we use only the traces corresponding to the internal

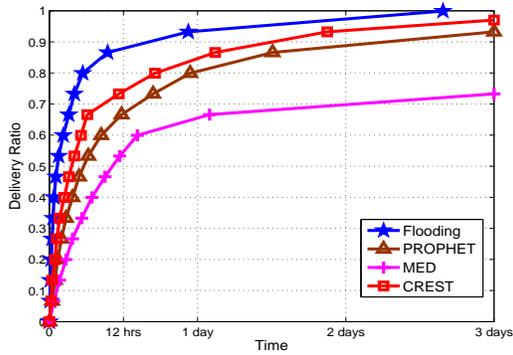
contacts. Although the office and conference datasets have several differences, we observe that a majority of the pairwise ICDs in the conference environment also follow log-normal distributions, which is consistent with previous findings [23].

Figure 5 compares CREST with the flooding, PROPHET, and MED schemes using the conference dataset for the single copy case. Figure 5(a) shows that the performance of CREST is marginally lower than that of flooding. In a 24-hour period, flooding delivers 95% of the messages, whereas CREST delivers 85%. Compared to PROPHET and MED, which deliver 80% and 60% of the messages in 24 hours, CREST has a lower end-to-end delay and a higher delivery ratio. Figure 5(b) compares the hop count for the messages delivered by the different schemes. The mean hop count for MED, flooding, CREST, and PROPHET is 2.28, 3.18, 4.15, and 6.72, respectively. These results are consistent with the results that we obtained using the mobility traces collected from the office environment. These results show that the CREST protocol can adapt and forward messages efficiently in different environments. We have also evaluated the performance of CREST by using mobility traces from other environments and synthetic traces in which the ICDs follow distributions that are not lognormal. CREST performs well in each case. Due to the page limit, we plan to present these additional results in an extension.

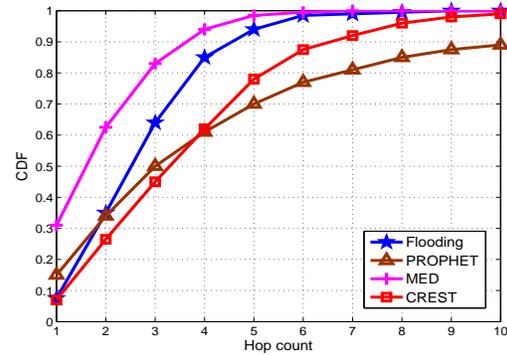
VI. CONCLUSIONS

In intermittently connected environments, it is advantageous to utilize the contact opportunities presented by mobile entities, such as humans and vehicles, to forward messages. However, it is important to judiciously choose the forwarding opportunities in order to reduce the end-to-end delay and provide a high delivery ratio, while incurring a low overhead. This can be better accomplished by characterizing realistic mobility traces and further developing practical forwarding protocols based on the models.

In this paper, we analytically characterized human mobility traces that we collected in our office environment, and observed that the inter-contact duration between pairs of nodes are not memoryless. Following this, we introduced a new link



(a) CDF of end-to-end delay



(b) CDF of the hop count

Fig. 5. Comparison of opportunistic forwarding protocols in a conference environment: single copy.

metric, median conditional residual time, which uses the previous contact time between a pair of nodes to accurately estimate the remaining time for their next contact opportunity. The median conditional residual time forms the basis of CREST, an opportunistic forwarding protocol that is completely decentralized and has minimal overhead. Experimental results show that CREST performs better than protocols that use future contact schedules and global knowledge of the contact behavior across the network. Furthermore, by disseminating only a few additional copies of the message at the source, CREST achieves a delivery ratio comparable to that of the flooding protocol.

The primary focus of our work so far has been to study the impact of human mobility on opportunistic forwarding. However, an office environment is also a social network, where people meet others based on their functional roles and the ties they have to different groups. The social structure of the network is less likely to change frequently compared to mobility-based contact traces. Metrics based on such long-term relationships provide an alternative way to represent contact behavior between pairs of nodes. Hence, it would be interesting to explore the use of novel metrics that characterize social ties and consider any correlations that may exist among the pairwise ICDs, in order to study their impact on opportunistic forwarding.

ACKNOWLEDGMENTS

We are grateful to the management at Deutsche Telekom Labs for supporting this project and the summer internship of the first author. We thank Olivier Dousse for his valuable feedback during the initial part of this work.

REFERENCES

- [1] X. Zhang, J. Kurose, B. Levine, D. Towsley, and H. Zhang, "Study of a Bus-based Disruption-Tolerant Network: Mobility Modeling and Impact on Routing," in *Proc. of Mobicom*, 2007.
- [2] J. Burgess, B. Gallagher, D. Jensen, and B. Levine, "MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks," in *Proc. of INFOCOM*, April 2006.
- [3] P. Hui, A. Chaintreau *et al.*, "Pocket Switched Networks and Human Mobility in Conference Environments," in *Proc. of Workshop on Delay Tolerant Networking*, August 2005.

- [4] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein, "Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet," in *Proc. of ASPLOS-X*, October 2002.
- [5] A. Pentland, R. Fletcher, and A. Hasson, "DakNet: Rethinking Connectivity in Developing Nations," *IEEE Computer*, vol. 37, no. 1, pp. 78–83, January 2004.
- [6] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke, "Data Collection, Storage, and Retrieval with an Underwater Sensor Network," in *Proc. of SenSys*, November 2005.
- [7] N. Nguyen, S. Krishnamurthy, P. Xie, and D. Jones, "Improving Information Availability in Storage-Centric Sensor Networks," in *Proc. of Local Computer Networks (LCN)*, October 2008.
- [8] 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. of INFOCOM*, 2006.
- [9] T. Karagiannis, J. Boudec, and M. Vojnović, "Power Law and Exponential Decay of Inter Contact Times Between Mobile Devices," in *Proc. of MobiCom*, 2007.
- [10] A. Vahdat and D. Becker, "Epidemic Routing for Partially Connected Ad-Hoc Networks," Duke University, Tech. Rep., 2000.
- [11] S. Jain, K. Fall, and R. Patra, "Routing in a Delay Tolerant Network," in *Proc. of ACM SIGCOMM*, 2004.
- [12] E. P. C. Jones, L. Li, J. K. Schmidtke, and P. A. S. Ward, "Practical Routing in Delay-Tolerant Networks," *IEEE Transactions on Mobile Computing*, vol. 6, no. 8, pp. 943–959, August 2007.
- [13] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic Routing in Intermittently Connected Networks," in *Proc. of SAPIR*, 2004.
- [14] J. Francois and G. Leduc, "Routing Based on Delivery Distributions in Predictable Disruption Tolerant Networks," *Ad-Hoc Networks*, vol. 7, no. 1, January 2009.
- [15] E. Daly and M. Haahr, "Social Network Analysis for Routing in Disconnected Delay-Tolerant MANETs," in *Proc. of MobiHoc*, 2007.
- [16] P. Hui, J. Crowcroft, and E. Yoneki, "BUBBLE Rap: Social-based Forwarding in Delay-Tolerant Networks," in *Proc. of MobiHoc*, 2008.
- [17] H. Cai and D. Eun, "Crossing Over the Bounded Domain: from Exponential to Power-Law Inter-Meeting Time in MANET," in *Proc. of MobiCom*, 2007.
- [18] Ekahau system, <http://www.ekahau.com/>.
- [19] L. Kleinrock, *Queueing Systems Volume I: Theory*. Wiley, 1975.
- [20] EasyFit: Distribution Fitting Software, <http://www.mathwave.com/>.
- [21] A. Papoulis and S. Pillai, *Probability, Random Variables and Stochastic Processes*. McGraw Hill, 2002.
- [22] A. Balasubramanian, B. Levine, and A. Venkataramani, "DTN Routing as a Resource Allocation Problem," in *Proc. of ACM Sigcomm*, 2007.
- [23] V. Conan, J. Leguay, and T. Friedman, "Characterizing Pairwise Inter-contact Patterns in Delay Tolerant Networks," in *Proc. of ACM Intl. Conference on Autonomic Computing and Communication*, 2007.